

**THE SOCIAL EPISTEMOLOGY
OF EXPERIMENTAL ECONOMICS**

DE SOCIALE EPISTEMOLOGIE VAN EXPERIMENTELE ECONOMIE

THE SOCIAL EPISTEMOLOGY OF EXPERIMENTAL ECONOMICS

DE SOCIALE EPISTEMOLOGIE VAN EXPERIMENTELE ECONOMIE

Thesis
to obtain the degree of Doctor from the
Erasmus University Rotterdam
by command of the
rector magnificus
Prof.dr. S.W.J. Lamberts
and in accordance with the decision of the Doctorate Board.

The public defence shall be held on
Thursday 9th of November 2006 at 13.30 hours

by

Ana Cristina Cordeiro dos Santos
born in Lisbon, Portugal.

ISBN 90 9021 076 8

© Ana C. Cordeiro dos Santos, 2006

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including information storage and retrieval) without permission in writing from the author.

Cover Design by Luis Carlos Amaro

See <http://eps.eur.nl/dissertations> for the electronic version of the thesis

Printed by Gráficos à Lapa, Lisbon, Portugal

DOCTORAL COMMITTEE

Promotor:

Prof.dr. I. U. Mäki

Other members:

Prof.dr. D. McCloskey

Prof.dr. J. Sonnemans

Dr. J.J. Vromen

**Para os meus pais,
Conceição e António**

CONTENTS

Acknowledgements	ix
1 Introduction: Epistemology, Experiments and Economics	1
Part One: The Social Epistemology of Experiment	
2 Experiment, Coherence and Materiality	25
3 The Argument from Sociality	61
4 The Social Epistemology of Experiment	89
Part Two: The Social Epistemology of Experimental Economics	
5 The Market Experiment and the Foundation of Experimental Economics	107
6 The Epistemological Arguments of Experimental Economics	137
7 Human agency (or lack thereof) in Economic Experiments	175
8 'Human Agency' and Inference	193
9 Preference Reversals or Good Experimental in Economics	209
10 Conclusion	237
References	243
Appendix 1	21
Appendix 2	57
Appendix 3	173
Summary (Dutch)	257
Curriculum Vitae	258

ACKNOWLEDGMENTS

The Social Epistemology of Experimental Economics (SEEE) is about the inescapably collective nature of knowledge production, in time and space. Even though this thesis is singly authored, it certainly would have not been possible without the support and collaboration of a great many. The first accountable for the SEEE are my colleagues from the Economics Department at the Instituto Superior de Ciências do Trabalho e da Empresa (ISCTE), Lisbon. Special thanks are due to Helena Lopes, José Maria Castro Caldas and Raul Lopes, who supported this project from the very beginning and helped me to create the conditions for its successful accomplishment. Helena and José have extra responsibility for they introduced me to the exciting field of the philosophy of economics as a fruitful area of research to help answering our common concerns and anxieties. I also remember the on-going discussions we managed to keep alive and which gave me an extra reason for wishing a more active participation in future collective endeavours. Ana Costa, João Rodrigues, Luís Francisco Carvalho and Ricardo Mamede, neither are you innocent here.

Having arrived at Rotterdam I found an exciting environment at the Erasmus Institute for Philosophy and Economics (EIPE). Inspiring teachers, provocative fellow students, and a host of visiting researchers who presented at the EIPE seminars first-hand accounts of the most recent contributions to the field. Arjo Klamer, Jack Vromen, Jonh Davis, John Gronewegen, Gregory Dow, Wade Hands, Mark Blaug, Deirdre McCloskey, Aki Lehtinen, Caterina Marchionni, Emrah Aydinonat, Frank Hindriks, Gülbahar Tezel, Roberta Muramatsu, you are, too, an important part of this work.

The more active contribution, of course, belongs to my supervisor, Uskali Mäki, from whom I learned how to produce professional academic work and learned to appreciate the merits of precision and rigour. However, perhaps more importantly, I learned from him the importance of sheer and shared enthusiasm in the business of science, whose outcomes are not always certain and predictable. Your enthusiasm was, at points, the most tangible thing I could hold on to. Part of this package is the amazing freedom you gave me to take the paths of my own choosing and the confidence you demonstrated in having done so. Thanks for all that.

There were also other direct contributions, specially the feedback I received from my written work. I gratefully acknowledge helpful comments from Igor Douven, Marteen Janssen, George Hendrikse, Emrah Aydinonat, Frank Hindriks, and Caterina Marchionni. I thank Aki Lehtinen who also helped me to improve the English of SEEE. A special thanks is due to Francesco Guala, who not only generously

commented on a substantial part of my thesis, but also provided hospitality during my research visit at the University of Exeter, UK.

Besides the intellectual support, I must also gratefully acknowledge the financial support of Fundação para a Ciência e a Tecnologia (Grant SFRH/BD/2822/2000), Portugal, and Vereniging Trustfonds Erasmus Universiteit Rotterdam, The Netherlands. FCT funded my visiting periods in Rotterdam. Vereniging Trustfonds funded the visit to the University of Exeter, UK, in 2005. Last but not the least, I must thank Juliano to whom I owe my endurance in the making of this thesis and five fantastic woman that became good friends and that have made this the most rewarding experience: Cate, Carlo, Gul, Mar and Rô.

CHAPTER 1

INTRODUCTION: EPISTEMOLOGY, EXPERIMENTS AND ECONOMICS

Man, being the server and interpreter of nature, can only do and understand so much ... as he has observed in fact or in thought the order of nature: beyond this he neither knows anything nor can he do anything. Neither the naked hand nor the understanding left to itself can effect much. It is by instruments and helps that the work is done, which are as much wanted for the understanding as for the hand. And as the instruments of the hand give motion or guide it, so the instruments of the mind supply either suggestions for the understanding or cautions. (Bacon 1859 [1620])

1.1 THE SOCIAL EPISTEMOLOGY OF EXPERIMENTAL ECONOMICS

The principal aim of this dissertation is to study the processes in which knowledge is produced and established in experimental economics. This investigation requires the construction of a framework by reference to which the experimental practice of economics can be comprehended and evaluated.

At first this might seem a rather unnecessary endeavour given that equivalent attempts might be expected to have already generated a comprehensive account of experimental practice that captures the underlying principles of knowledge production and justification. After all experimentation has been a constitutive part of knowledge production since as early as the seventeenth century. However, it was not until the latter quarter of the twentieth century that experimentation attracted the interest of philosophers of science. This is not to say that the role of experiments in the progress of science was denied. It simply means that until then experiments were not considered a philosophically interesting subject matter.

In spite of the views of a couple of visionary philosophers, like Francis Bacon, who soon realised the significance of the interactive work performed by the ‘hand’ and the ‘mind’, most philosophers in the past continued to focus on the “the understanding left to itself”. As a result, theory continued to be perceived of as the complete expression of scientific accomplishment and the proper subject matter of the philosophy of science.

The interest in scientific experimentation was motivated by the rise of the sociological studies of science, which denied experiments the adjudicator role traditionally ascribed to them, either because theories are immune to empirical evidence or because experimental reports fail to provide neutral evidence of nature.

It is therefore not surprising that philosophical studies of experimentation are still very fragmented. A comprehensive account of experimental practice that captures the

underlying principles of knowledge production and justification is not yet available. Constructing such a framework is thus the first goal of the present research. The second goal consists of applying this framework to the social processes of knowledge production, the method, the arguments, the procedures, the designs and the results of experimental economics.

The framework for the analysis of scientific experimentation put forward in this thesis – *The Social Epistemology of Experiment* – aims to provide an analytical device that enables the identification of the basic attributes of scientific experimentation that justify experimenters' beliefs in knowledge generated by laboratory experiments. The epistemology of experiment is a *social* epistemology because it emphasises the epistemic role of the social dimension of the experimental process of knowledge production.

The objective of this thesis is not only to paint, in broad strokes, the field of experimental economics and the subfields of research that make it up. In addition, it also aims to provide an accurate and detailed depiction of particular research programs, single experiments, experimental techniques and experimental results. This thesis therefore attempts to identify the defining characteristics of the experimental method of economics, its standard practices, justifying arguments and central debates, along with the careful scrutiny of particular experiments, procedures and arguments, which taken together contribute to a simultaneously comprehensive and precise portrait of experimental economics. However, before presenting the *Social Epistemology of Experiment*, a brief introduction of the study of experiments is required.

1.2 THE NEGLECT OF EXPERIMENTS IN THE STUDIES OF SCIENCE¹

Experiments have been part of the scientific process of knowledge production since as early as the seventeenth century. We are all familiar with mythical stories of revolutionary experiments that have changed western worldviews – think of Galileo's discoveries with the telescope or Newton's investigations of the refraction of light with the glass prism. However, it was not until the latter quarter of the twentieth century that experiments attracted the attention of philosophers. The empirical evidence experiments supplied was acknowledged as important. Nonetheless, the scientists' practical engagement with the natural and social worlds to produce them was not considered an interesting subject matter. Curiously, perhaps, scientific experimentation only captured the attention of the philosophers when the status of empirical evidence began to be under attack.

There is no doubt that since the dawn of scientific experimentation, philosophers addressed the role that experiments play in science in some way or another. Certainly

philosopher-scientists such as Galileo Galilei (1564-1642) and Isaac Newton (1642-1727) did not ignore them. Moreover, already at the end of the sixteenth century, Francis Bacon (1561-1626) noted the importance of the interactive work performed by the 'hand' and the 'mind' and the sterility of the 'naked hand' and that of the 'understanding left to itself'. As Bacon remarked, the 'naked hands' of the practical scientists (who resemble ants) can only collect and use their gatherings whereas the 'minds' of the theoretical scientists (who resemble spiders) can only make cobwebs out of their own substance. Scientists should be instead like the bee "that gathers material from the flowers of the garden and the field, but transforms and digests it by a power of their own" (1859 [1620]). Indeed, the course of science followed Bacon's prescription and it did so at an ever-increasing pace. At present, the interactive work performed by the hand and the mind encompasses, and gives rise to, new fields of research of which experimental economics is only one among various examples. Nonetheless, in the philosophy of science not only do theory and experimentation continue to be relatively distant from each other, but theory also continues to be perceived of as the complete expression of scientific accomplishment and the proper subject matter of philosophy.

The causes of the neglect of experiments can perhaps be traced back to the classical divide between Rationalism and Empiricism that separated 'reason' and 'experience' as the two opposite foundational bases of scientific knowledge.² This separation strongly contributed to keeping apart the 'theoretical' and the 'observational' aspects of knowledge production. However, the Empiricists' neglect of work done 'by the hand' is perhaps more surprising. Thus, while rationalism overlooked experiments by privileging the *a priori*, self-evident and certain knowledge, empiricism neglected experiments by equating empirical evidence to knowledge acquired by the immediate, passive and unmediated senses. In the nineteenth century, empiricist philosophies took over but experience continued to be perceived as unproblematic sense perception. Of course, there are some notable exceptions such as that of the empiricist John Stuart Mill (1843) who, in his *System of Logic*, provided the first elaborated account of experimentation as the method of difference whereby cause-effect relations can be identified.

In the twentieth century, with logical positivism, 'reason' and 'experience' became both acceptable sources of knowledge in philosophical writings.³ However, they continued to belong to distinct realms. While the former was *a priori* knowledge made of the formal and logically true analytical propositions of logic and mathematics, the latter was composed of the factually true synthetic propositions of science. Moreover, experience remained the foundational basis of knowledge.

² To name a few, prominent rationalists are the Continental philosophers René Descartes (1596-1650), Benedict de Spinoza (1632-77) and Gottfried Wilhelm Leibniz (1646-1716) and the empiricists include the British philosophers John Locke (1632-1704), George Berkeley (1685-1753) and also David Hume (1711-76).

³ This view is attributed to a group of thinkers who came to form, in the 1920s, what would be known as the Vienna Circle. The Vienna Circle included the names of Moritz Schlick (1882-1936), Rudolph Carnap (1891-1970) and Otto Neurath (1882-1945).

¹ This section briefly sketches the conception of experiment in the philosophy of science to set the background for the present research. Because this history has already been competently told by others (Hacking 1983; Giere 1988; Mayo 1996) only the main aspects of the neglect of scientific experimentation and the principal reasons for its recent interest are highlighted.

Theories, to be regarded as scientific, had to be constructed inductively from sentences about empirical phenomena (i.e. the ‘protocol language’) or they had to be translatable into them by ‘correspondence rules’.⁴ Experiments were simply an unproblematic source of observational reports.

The sharp contrast between the theoretical and the observational later became subjected to criticism, in the second half of the twentieth century, under logical empiricism.⁵ Along with this tradition came a more informed inclusion of experiments in philosophical writing. The relation between theory and evidence was no longer conceived of as a matter of building up generalisations from observational statements or of translating theoretical statements into an observational language. Rather, it became instead a matter of submitting the predictive consequences of the theory to empirical testing. Following the hypothetico-deductive method,⁶ only the predictive statements, deduced from a general hypothesis and statements of initial conditions, were to be subjected to empirical testing. The inversion of the relation between theory and observation made it evident that empirical testing requires special conditions to test theoretical hypotheses. In particular, it requires the construction of specific test systems that satisfy the set of initial conditions assumed as true in the derivation of the theory’s predictions. The explicit recognition that empirical testing involves the construction of a specific test system opened up the debate about the status of “observational statements” that ultimately led to the study of experiments.

Different views about the status of “observational statements” were presented or became relevant. Karl Popper (1959[1934], 1965) explicitly recognised that ‘observation statements’ were not neutral claims about the world. Observation statements are always ‘interpretations in the light of theories’. They are also ‘fallible’ because there is no way to ascertain their true value. Consequently, their acceptance relies on agreement by ‘convention’ given the set of ‘background knowledge’. The fragile status of observational statements, in turn, undermined the adjudicator role of empirical evidence. Nonetheless, the philosophers continued to perceive the bulk of scientific work as consisting of the derivation of theoretical predictions. They continued to think of scientific experimentation as an activity subservient to and dependent on theory:

The theoretician puts certain definite question to the experimenter, and the latter by his experiments tries to elicit a decisive answer to these questions, and to no others. All other questions he tries hard to exclude... it is a mistake to suppose that the experimenter... [aims] ‘to lighten up the task of the theoretician’, or... to furnish the theoretician with a basis for inductive generalisations. On the contrary, the theoretician must long before have done his work, or at least the most important part of his work: he must have formulated his questions as sharply as possible. Thus, it is he who shows the experimenter the way. But even the experimenter is not in the main engaged in making exact observations; his work is largely of a theoretical kind. Theory dominates the experimental work from its initial planning up to the finishing touches in the laboratory. (Popper 1959 [1934], p. 107)

The *theory-ladenness* of observation reports has been a well-known idea since the 1960’s as testified by many important contributions including those of Michael Polanyi (1958), Norwood Russell Hanson (1958), Thomas Kuhn (1970[1962]), Imre Lakatos (1970) and Paul Feyerabend (1975). Another important critique is conveyed by the *underdetermination thesis* popularised by Willard Van Orman Quine (1953), but also credited to Pierre Duhem (1954[1906]). This attacked the purportedly univocal relation between the theoretical hypothesis under test and the empirical observations. According to this thesis, any empirical test is a conjoint test of a target hypothesis together with a set of auxiliary hypotheses about the test system which renders virtually impossible the attribution of a failure to the accountable element(s). The implication of this is that an inconsistency between theory and evidence may always be attributed to some auxiliary assumption rather than the hypothesis under test. More generally, the scientific worldview is part of a ‘web of belief’ (Quine 1980) of mutually intertwined scientific beliefs that are revised with experience. The decision of what parts of the web to revise depends on the context and on pragmatic considerations rather than on an explicit application of a particular piece of evidence to a particular belief. From this it follows that conflicts between theories, theoretical hypotheses or beliefs with experience result in human decisions concerning which part of the theoretical and/or empirical system is to be modified.

It was the historical approach to science that emerged in the 1960s that ultimately assailed the epistemic power of observation. It did that by showing that science does not follow the pattern described by the proposal, the testing and the rejection of theories, followed by the subsequent proposal of novel theories. Thomas Kuhn’s *Structure of Scientific Revolutions* (1970[1962]) was most influential in challenging this position and replacing it with the view that scientists are instead engaged in minor theoretical tinkering within the limits defined by a ‘paradigm’ that is not questioned. Moreover, when one paradigm is replaced by another it is not followed by critical discussion based on empirical and reasoned arguments. Rather, its replacement is determined by contingent social conditions given that there is no objective basis, empirical or otherwise, on which to base a paradigmatic change. Paradigms are ‘incommensurable’. They convey different ‘worldviews’ and are characterised by different methods, theories, instruments and standards.

Even though the various practical aspects of conducting an empirical test began to emerge in philosophical discussions, the issues of the theory-ladenness of observation, the underdetermination of theory by data, and paradigm incommensurability were still very much informed by a theory-oriented view of science. They all pointed to the absence of a neutral empirical basis for theory testing and choice. The actual processes of knowledge production, and in particular those concerned with the construction of the empirical basis of science, remained largely ignored. However, these criticisms stimulated an interest in determining how the empirical basis of science is actually arrived at. This recent history of the philosophy of experiment is eloquently summarised by Robert Ackerman in the following paragraph:

The philosophy of science that developed out of 20th century positivism placed a heavy foundational emphasis on observational fact as the means of controlling theoretical growth,

⁴ Cf. Carnap (1967[1928]).

⁵ Important contributors to logical empiricism were Richard Braithwaite, Carl Hempel and Ernest Nagel.

⁶ Cf. Hempel and Oppenheim (1948).

but although theoretical statements were logically articulated against observational statements in increasingly sophisticated ways as positivism developed, positivism paid little attention to the way in which statements of observational fact were produced in experimental practice. One simply began to philosophise on the assumption that science was capable of delivering a data base of settled observational statements. In the vengeful dismantling of positivism undertaken after Kuhn's work, the old connotations of *fact* were replaced by assertions that theoretical expectations somehow determined the observations of science, and later by suggestions that observations are constructed by groups of scientists engaged in a social process of negotiation (1989, p. 185 emphasis in original)

The study of the actual processes of production of experimental knowledge was first carried out by sociologists. The so-called 'strong programme', or the Edinburgh School, associated with the work of the sociologists Barry Barnes (1974) and David Bloor (1976), took the process of belief formation in science as part of the study of sociology, as the processes of belief formation by other social groups already were. According to the new branch of the sociology of scientific knowledge (SSK), scientific beliefs are to be explained by the social interests of the scientists, which are determined by the position scientists occupy in the overall pattern of social relationships. The important contribution of SSK, which is of relevance here, was that experiment became a subject matter of inquiry in its own right. Detailed case studies were carried out from the historiography of science and from sites of scientific practice (based on field work, ethnographic inquiry and participant observation) that revealed how experimental knowledge is actually arrived at. The long sequential phases of trial-and-error, the dead-alleys, the failures involving the construction and calibration of material apparatuses, the importance of experimenters' personal and social skills, the historical and social milieu embedding scientific production, all became part of the written accounts of experiments. The view that experimentation is an unproblematic activity that follows a linear and logical sequence of actions leading to a conclusive result was no longer part of the stories of experiment.⁷

Two accounts of experimental practice had a tremendous impact on subsequent work on experiment: Bruno Latour and Steve Woolgar's *Laboratory Life* (1986[1979]) and Harry Collins's *Changing Order* (1985). The former, adopting the anthropological approach to a community of scientists, focused on the material, technical, literary and social practices undertaken to transform initial 'literary inscriptions' generated by 'inscription devices' into scientific papers, allegedly the main goal of scientific activity. The epistemological import of this study was that the acceptance of experimental results as 'scientific facts', involves a long process of social negotiation taking place both in the laboratory and outside. Collins focused on the lack of scientific standards for the assessment of experiments that produce novel phenomena. The problem, as he saw it, is that when a clash with an expectation occurs, scientists must decide whether to

trust that the apparatus has worked properly and thereby accept the experimental result, or to attribute the result to a failure of the apparatus, thus rejecting the experimental result. In the absence of a criterion on the basis of which the proper functioning of the apparatus can be appraised, there is no secure way to ascertain what the proper outcome is. Conversely, because there is no available account of the phenomenon, neither are there means to ascertain whether the apparatus operated adequately. This is how Collins presented the problem of the 'experimenter's regress' in the study of gravity waves:

What the correct outcome is depends upon whether there are gravity waves hitting the Earth in detectable fluxes. To find this out we won't know if we have built a good detector until we have tried it and obtained the correct outcome! But we don't know what the correct outcome is until... and so on *ad infinitum*. (1985, p. 84)

From this perspective, the absence of a natural stopping point in experimentation leaves room for a high degree of arbitrariness, which renders experimental results vulnerable to scientists' predispositions and to the interference of non-epistemic factors. As Collins states: "[i]t is not the regularity of the world that imposes itself on our senses but the regularity of our institutionalized beliefs that imposes on the world" (1985, p. 148).

Moreover, the private nature of the 'tacit knowledge' (cf. Polanyi 1958) involved in the knowledge production process renders the experimental process of knowledge production relatively inaccessible to those who do not take part in it. This means that not only are there no *a priori* criteria to assess experimental results, but also there are no intersubjective criteria to that end. The implication of this is that other scientists are excluded, if not prohibited, from the external validation of experimental results. The difficulty of replicating experimental results generated by others (i.e. the difficulty of reproducing somebody else's experimental process) renders the decision of what counts as a replication also contextual and open to negotiation. Furthermore, there is little incentive to actually attempt to replicate somebody else's results. On the one hand, the reproduction of pretty consensual results is not very stimulating to scientists. On the other hand, if the result of the experiment is considered invalid, it makes no sense to try to reproduce it. Only highly controversial results are worth attempting to reproduce.

In summary, the problems presented by theory-ladenness, the underdetermination thesis, incommensurability, the experimenter's regress and the tacit nature of experimental knowledge, opened up experiments to the interference of 'social factors' in knowledge production. The lack of universal criteria for the appraisal of experimental results, as well as the private nature of knowledge production, seriously undermined the enterprise of experimentation. Thus, from an unproblematic activity that generated knowledge about the natural world, scientific experimentation became an activity prone to the influence of factors external to science. Given that experimentation had long been an activity that was taken for granted, it is not surprising that the interest in experiment arose when its privileged status was called into question. Philosophers became the next students of experiments.

⁷ The exclusion of the details of experimental practice from the published reports of experiments explains, at least partially, the neglect of experimentation by the students of science. These reports present highly idealised accounts of experiments where experimental practice is depicted as a fairly unproblematic activity whose results have direct bearing on some theoretically interesting question.

1.3 THE NEW EXPERIMENTALISM

The view that scientific practice must be understood within its cultural milieu is a position that most philosophers were willing to accept. However, the view that the content of science is also prone to the influence of factors external to science (e.g. social prejudices, ideologies, etc.) was too radical to be left unnoticed. By questioning science, the sociological studies were not only targeting science, but they were also attacking the basis of the ‘rationality’ or ‘objectivity’ of science. In this situation, more conservative historians and philosophers attempted to overcome the problems of the traditional view (i.e. the unrealistic depiction of science and the unattainable prescriptions for good scientific practice) while avoiding the radical implications of SSK (i.e. the irrationality of science). They did that by having recourse to the actual practice of science to demonstrate that experimental practice is not irrational and that the experimental results are not very vulnerable to extra-scientific influences. The more moderate sociologists, perhaps concerned with the implications of the application of SSK onto itself, also attempted to reduce the impact of the ‘social’ by highlighting the role of ‘nature’ in knowledge production. The point was that nature cannot be shaped at will no matter how strong scientists’ beliefs and desires are.

In this context, a new philosophical approach emerged within which the credibility of experiments, and hence of science, would have to be founded on actual experimental practice, i.e. on the strategies, instruments, techniques, actions, and decisions that ground belief in experimental results. This is how Ackermann perceived the contribution of the new research programme:

[A] number of recent studies have gone back to the history of science to study experimentation as a means of grounding the nature and origins of the observational facts that many still think must offer important objective constraints for scientific theory, at least at certain pivotal points in scientific development. Empiricism is in this way resurrected (although transformed) by finding a philosophical account of experimental data that are especially worth having, and that function as solid hinge points for controlling theoretical conjectures...

[T]he newer concentration on experimentation is that an experiment is a complex activity undertaken over time (involving the design and manufacture of equipment, the calibration of equipment, checks on proper functioning of the equipment, etc.) that may issue in observations that can be reported as data. *What’s needed in this context is a discussion of whether specific experimental practices can in some sense legitimate or validate observational reports, and how the strength of such legitimation might be taken into account in a philosophy of science.* (1989, p. 186 emphasis added).

New detailed accounts of laboratory science have since been carried out to identify the constitutive aspects of experimental practice and reasoning that could restore to experimentation its traditional function as an objective adjudicator in science. The goal of the new experimental approach is stated by Peter Galison in the following way:

It is unfair to look to experimental arguments for ironclad implications and then, upon finding that experiments do not have logically impelled conclusions, to ascribe the experimentalists’ beliefs entirely to “interests”. But who would have thought that experiments were like mathematics? *It seems to be more promising to look at how experimentalists do reason, build and test their way to the assembly of persuasive evidence*, even though the experimentalists’ reasons for belief lack the form of a mathematical demonstration. (1987, p. 11, emphasis added)

And in the words of Allan Franklin:

In these postmodern times, when scholars proclaim the impossibility of knowledge, except, of course, their own, I would like to defend a rather old-fashioned notion. This is the idea that science provides us with knowledge about the world which is based on experimental evidence and on reasoned and critical discussion. I further suggest that we can, and we do, have reasonable grounds for belief in that evidence. In short, I believe that science is a reasonable enterprise. (1999, p. 1)

The first ‘new experimentalist’ contribution is credited to Ian Hacking’s *Representing and Intervening*, and specifically, to the second part of the book which concerns the intervening aspect of the title. The main message is that scientific experimentation is “a long hard task” whose goal is “to create, produce, refine and stabilise phenomena”, which calls for a wide range of activities not encompassed in the notion of passive observing. For “[o]nly when one has got the equipment right is one in a position to make and record observations” (1983, p. 230). The central question of experimentation when producing novel phenomena is how to determine when the experimental equipment is working correctly so that the observation of the phenomenon can take place, or how to be sure that what is observed is a real property of the natural world rather than an outcome of an incorrect performance of the experimental apparatus, i.e. an artefact. According to the new experimentalists the answer, or answers, must be found in experimental practice.

Hacking identifies two arguments to justify that a phenomenon produced is not the outcome of a malfunction of the experimental apparatus. The ‘*no miracle argument*’ points to the confidence provided by the production of the same phenomena by different material apparatuses. The upshot is that it is very unlikely that the phenomenon is a malfunction of different physical processes all generating the same ‘artefact’. It is instead more likely that the phenomenon observed is the outcome of the natural attributes of the object under observation. This argument is coupled with the *argument from intervention* that asserts that the consistency between the results of intervening with the material world, with what is known about it, grounds belief in the observations (Hacking 1983, Ch. 11).

Another set of arguments has been provided that intends to demonstrate the relative autonomy of experiments from theory. First, Hacking distinguishes different kinds of theoretical hypotheses which put the relation between theory and experiment into a new perspective, thereby undermining the theory-ladenness of observations. The point is that the theoretical presuppositions underlying an experimental result are low-level hypotheses (about the functioning of the apparatus, the adequate means of

operating with it, how the results are to be brought to bear on the phenomenon, and so forth) rather than a high-level theory that explains the phenomenon observed. Even though the experimental evidence may be theory-laden in some broad sense, it does not invalidate its role as empirical evidence for the test of theories which are not implicated in the production of experimental data. Second, it is noted that not all experiments are tests of high-level theories. Some experiments have ‘a life of their own’ in that they are implemented to refine instruments, obtain a sharper manifestation of a badly-understood phenomenon or to explore new domains (ibid, Ch. 9). Finally, Hacking notes the resilience of experimental knowledge to high-level theoretical change. Changes in theory are not necessarily accompanied by changes in experimental knowledge. In fact, the history of science demonstrates the stability of experimental knowledge, in the form of instruments, techniques, strategies, and results, which has been crucial to experimental activity and to the growth of experimental knowledge. It is the stability of instruments and the knowledge of the properties of well-known entities that allows scientists to investigate the less-well-known aspects of nature (ibid, Ch. 12). In summary, notwithstanding the fact that experimental statements are neither theory-free nor infallible, the relative autonomy of experiments from theory and the relative stability of experimental knowledge justify experimentation having its traditional role as a neutral adjudicator.

The last two decades has been prolific in the production of detailed accounts of scientific experimentation in the natural sciences. The most influential authors include Ian Hacking, Allan Franklin, Peter Galison, Andrew Pickering, and David Gooding. Volumes of collective papers have been edited by Achinstein and Hannaway (1985), Gooding, Pinch and Shaffer (1989), Pickering (1992), Buchwald (1995), Heidelberger and Steinle (1998) and Radder (2003). And several journals, for the first time, devoted special issues to the philosophical discussion of experiment (e.g. the volumes of the symposia of Philosophy of Science Association 1988, 1990, 1996). Surveys of this literature have also been made available and can be found in Ackermann (1989), Hacking (1989), Pickering (1992), Mayo (1996), and Franklin (1998).

Because these studies have been intended as responses to particular charges informed by a theory-inspired view of science (e.g. theory-ladenness, underdetermination, the social conditioning of science, etc.) and to the ensuing debate that followed, they are partial and scattered accounts of experimentation. As yet there is no account that identifies the general principles guiding scientific experimentation that also explains what is it that is accomplished that grounds belief in experimentally generated knowledge. This is explicitly recognised in the most recent edited book in the field published two decades after Hacking’s seminal contribution:

The development of the philosophy of scientific experimentation over the past twenty years has two main features. After a rapid start in the 1980s... it seems to have lost much of this momentum during the next decade. At the very least, the expectation that the study of experiment would become a major issue within received traditions in philosophy has not been fulfilled. To verify this, it is enough to glance through the recent volumes of well-known journals, such as *Philosophy of Science*, *British Journal for the Philosophy of Science*, *Erkenntnis*, and the like. Alternatively, one may look at recent anthologies, which could be supposed to represent

the core readings in present-day philosophy of science. For example, the six-volume set of collected papers in philosophy of science (Sklar, 2000) contains no contributions that focus on experimentation. And in the voluminous *Companion to Philosophy of Science* (Newton-Smith, 2000) the explicit analysis of experimentation is almost completely limited to one chapter. Thus, the fact that many scientists, perhaps even the majority of them, spend most of their time doing experiments of various kinds is not reflected in the basic literature in the philosophy of science... thus, *the philosophy of experimentation is still underdeveloped, especially as compared to historical and social scientific approaches...* Given this state of affairs, many philosophers of experiment agree that the field needs a new impulse. (Radder 2003a, p. 1-2 emphasis mine)

The most elaborate and consensual account of experimentation to date sees experimentation as an activity that consists of the progressive elimination, measurement and calculation of ‘background effects’, ‘noise’, or ‘error’, until experimenters are confident that the phenomenon created in the laboratory was successfully isolated from disturbing factors (Franklin 1986, 1990a, 1999; Galison 1987; Hon 1989, 2003; Mayo 1996; Guala 2005b). Accordingly, as errors are eliminated, corrected and accounted for, belief is reinforced as to the ‘reality’ of the ‘signal’ looked for. The process stops when experimenters are convinced that they have accounted for all sources of error. However, as it will be argued, the conception of experiments as error-elimination endeavours does not illuminate what it is that experimental practice achieves that justifies confidence in experimental results. Nor does it identify the source of the epistemic value of the experimental process of knowledge production. In short, error elimination is a negative epistemology that highlights the successive eradication of known or identifiable sources of error during the process of knowledge production. The positive sources of epistemic value of the experimental processes and products are left out. Furthermore, the error elimination strategy provides a counterintuitive explanation of the formation of belief. Given that errors are identified by clashes between scientists’ expectations and the experimental outcomes, experimenters have justified belief whenever they succeed in confirming prior expectations because then they will believe that all sources of error have been eliminated. However, the research carried out here demonstrates that the confirmation of scientists’ anticipations does not suffice to ground belief in experimental knowledge. Additional arguments must be provided as to why scientists come to believe that all relevant errors have been eliminated. The Social Epistemology of Experiment provides such arguments by way of a positive account of scientific experimentation.

1.4 THE SOCIAL EPISTEMOLOGY OF EXPERIMENT

The *Social Epistemology of Experiment* (SEE hereafter) shares with *New Experimentalism* the principle that the methodological and epistemological study of experiment must take into account the actual processes by which experimental knowledge is produced. However, the former attempts to go further. It offers a general account of experiments that integrates the existing contributions into a unifying conceptual framework highlighting the fundamental attributes of scientific

experimentation. SEE presents a comprehensive account of the inner processes of knowledge production that is applied to the fundamental methodological and epistemological questions of experimental practice. Specifically, SEE identifies the basic attributes of scientific experimentation that justify the experimental process of knowledge production as a reliable process and the experimental results as justified beliefs. In short, SEE aims to answer the broad question:

What are the defining features of the experimental process of knowledge production that justify experimenters' belief in experimentally generated knowledge?

THE STRATEGY

SEE is built upon the new experimentalist accounts of laboratory practice in physics as well as the contributions of the recently created branch of philosophy called social epistemology (see Schmitt, 1994). The former encompass the reports of experimental physics by the philosopher Ian Hacking, the sociologist Andrew Pickering, and the historians Peter Galison, David Gooding and Allan Franklin. The latter includes, most significantly, the work of the philosophers Philip Kitcher, Miriam Solomon and Helen Longino. These sources are, however, complemented with other contributions to the study of experiments.

It should be noted at the outset that the prominent role of experimental physics in the construction of SEE is purely due to pragmatic considerations. Experimental physics has been the most studied experimental field from a variety of disciplinary viewpoints – philosophical, historical, sociological, and anthropological. Moreover, these studies draw epistemological implications. This is precisely the feature common to the selected studies. They all provide detailed historical accounts of the practice of experimental physics with implications for the epistemological issues this thesis addresses. Four remarks are relevant in this respect.

First, the studies are taken as accurate descriptions of experimental practice based on sources other than the final written accounts of experiments. They rely on various outputs of scientific practice, such as scientists' notebooks, private correspondence, visits to experimental sites, interviews with relevant scientists, participant observation, and so forth. Therefore, these studies are able to uncover the human action hidden in the reconstructed accounts of experiments relevant to the production and establishment of experimental knowledge. Second, all the studies attempt to provide arguments for the reliability (objectivity or rationality) of the experimental process of knowledge production that was shaken with the breakdown of the traditional view of science. Third, these accounts are representative of the diversity of the available contributions. They have in fact been part of arguments of different philosophical positions. Whereas some aim at restoring the traditional role of experiments (e.g. Franklin), or rescuing the sociological approach to science (e.g. Gooding), others can be placed in the continuum delimited by the more sociological-oriented and the more traditional conception of experiment (e.g. Galison). Fourth, the studies focus on different genres of experimental practice, with respect to time, place, size, and instruments. The contributions from social epistemology are used to elaborate on the

epistemic role of the social dimension of knowledge production that exists within the studies of experimental practice but have not been given sufficient attention.

THE EPISTEMIC ARGUMENTS OF EXPERIMENTS

SEE is presented in the form of three epistemic arguments that justify experimenters' belief in experimental knowledge and four epistemic tests for the appraisal of the experimental practice and products. The analysis of the studies of experiments concurs with an account of experimental practice as a scientific endeavour that aims to establish relations of coherence among heterogeneous items of scientific culture. Drawing very closely on Andrew Pickering's (1989, 1995b) account, the goal of experiments is reconstructed as the attainment of a *three-way coherence* among the three main components of the experimental system: the 'material procedure', the 'instrumental model' and the 'phenomenal model', or that between experimental practice, the conceptualisation of that practice and the conceptualisation of its outcome. The studies of experimentation suggest an *argument from coherence* that justifies experimenters' belief in experimental knowledge. It asserts that the three-way coherence achieved between the component parts of the experimental system provides scientists with reasons for believing in the experimental results. The alignment achieved between the components of the experimental system supports each single component and the system they make up and thereby the experimental result conveyed by the phenomenal model. Scientists then believe that they have produced the phenomenon of interest or they believe that they have controlled the influence of extraneous factors so that there is no need to introduce further amendments to the experimental system.

The simple attainment of a three-way coherence, however, does not provide sufficient justification for the experimental result it supports. The reason is that the epistemic value of the three-way coherences may be negatively affected by present scientific culture and scientists' prior commitments. On the one hand, scientific culture may constrain the generation of research problems with a high potential for the creation of novel knowledge. It does that by selecting problems that can be solved with the available conceptual and instrumental tools. On the other hand, scientists' prior commitments may limit the epistemic value of the solutions arrived at insofar as these are prone to scientists' predisposition to actualise favourite solutions. This predisposition may lead to the premature conclusion of the processes of knowledge production. The history of science shows that it is not rare for experimenters to prematurely put an end to experiments when expectations are met, while relevant courses of action are overlooked or neglected. To put it differently, an epistemically significant three-way coherence is one that is built to answer a meaningful scientific question following the exploration of relevant courses of action.

The argument from materiality and the argument from sociality provide further justification for the epistemic value of the three-way coherences. The *argument from materiality* asserts that the three-way coherence that sustains an experimental result endows the result with epistemic value given that the participation of the 'material world' in knowledge production promotes the revision of scientific culture and scientists' prior beliefs and thereby promotes the production of novel and reliable

claims to knowledge. The epistemic role of the 'material world' in knowledge production consists of its capacity to frustrate scientists' expectations which stems from scientists' inability to control the material world so as to make it conform to their expectations. The clash with expectations creates unforeseen problems, the resolution of which entails a high potential for the generation of novel knowledge. Nonetheless, the participation of the 'material world' may still be severely constrained by the plasticity of the experimental systems. Plastic systems can easily accommodate material resistances to both established beliefs and scientists' expectations.

The *argument from sociality* asserts that the three-way coherence that sustains an experimental result endows the result with added epistemic value because the collective aspect of knowledge production promotes the examination of scientific culture and scientists' expectations. The fact that knowledge production and justification are collective endeavours renders the results of science less vulnerable to the constraining effects of established beliefs and expectations. The confrontation of different cultures and commitments give rise to unforeseen problems that extend the process of knowledge production to the pursuit of courses of action taken to be relevant to the problem at hand. The implication of this is that the claims to knowledge that come to be accepted by the scientific community are epistemically valid because they carry novel knowledge which is supported by the articulation of various and varied items of scientific culture and beliefs.

The argument from sociality together with the argument from coherence and the argument from materiality provide a complete account of how experimenters generate and establish experimental knowledge. The three arguments in conjunction yield the following result:

The greater the number and the wider the heterogeneity of the resources (material, conceptual and social) involved in knowledge production, the higher the epistemic status of the three-way coherence and the experimental result it supports. The solution arrived at is epistemically sound because it is the result of practices that have explored relevant courses of action to finding a solution to the problem at hand.

Of course, heterogeneity is neither a sufficient nor a necessary condition for the production of epistemically sound results. Nonetheless, having recourse to a number of relevant heterogeneous resources is epistemically significant when experimenting with new experimental systems to produce novel phenomena.

THE FOUR EPISTEMOLOGICAL TESTS OF EXPERIMENTATION

The epistemic arguments of experimentation suggest four *epistemic tests* for the appraisal of experimental claims to knowledge: 1) the materiality test; 2) the stringency test; 3) the social robustness test; and 4) the technological test.

The *materiality test* assesses the *material quality* and the *material agency* of the experimental system, i.e. it examines the composition of the experimental system and the actual participation of the material world in knowledge production. The epistemic import of the materiality test is the following: the higher the material

quality of the experimental system and the higher the agency of the material world, the higher is the epistemic value of the experimental result.

The *stringency test* assesses the rigidity of the experimental system, i.e. it evaluates the extent to which an experimental result may be rendered compatible with an item of scientific culture, or an expectation, by material or conceptual manipulation of some part of the experimental system. The epistemic import of the stringency test is the following: the higher the degree of stringency of the experimental system, the higher is the epistemic value of the experimental result.

The *social robustness test* evaluates the social process that gave rise to the experimental result. It calls for judgement about the critical nature of the social context that produced the experimental claim to knowledge. That is, it assesses the extent to which the experimental process of knowledge production answered relevant questions to the problem at hand and the experimental result is supported by varied and heterogeneous items of scientific culture. The epistemic import of the social robustness test is the following: the higher the social robustness of the knowledge production process the higher is the epistemic value of the experimental results.

The *technological test* assesses the degree of applicability of the products of scientific experimentation to the solution of new problems. This is a test of the subsequent contribution of the results of science to scientific and non-scientific purposes alike. This is a social requirement of science that demands that the products of science should be useful and contribute to the betterment of human life.

THE SPECIFICITY OF THE SOCIAL EPISTEMOLOGY OF EXPERIMENTS

SEE highlights two interrelated modes of analysis of experimentation that have not yet been fully articulated in epistemological studies of scientific experimentation. First, it integrates the analysis of the processes of production of experimental knowledge with the analysis of the products of experiments. The point being made is that the epistemic analysis of a phenomenon, instrument, procedure, and so forth, is inseparable from the analysis of the processes in which those experimental products came about. Second, the analysis of the processes in which experimental products are produced and established requires the appraisal of the participation of both the material and the social worlds.

The studies of experimentation undertaken to date analysed the experimental practices and arguments that help to justify the reliability of the experimental production process and products. However, as will be shown, they do not make it explicit how and the extent to which, at a given moment of practice, the reliability of the process of knowledge production may be variable and thus bestow differentiated epistemic value on the experimental results. The studies are still very much attached to a finished picture of science that focuses on well-established results for which sufficient supporting arguments have already been gathered.

The 'materiality' and the 'sociality' of scientific experimentation have been recognised as relevant components of scientific experimentation. However, they have not been considered together. In fact, they are often presented as epistemically opposite and contrasting factors of knowledge production. On the one hand, the

material component of experimental systems is generally presented as the countervailing factor to the detrimental effect of the social factors in knowledge production, while the positive effects of the latter are overlooked. On the other hand, the evocation of the social dimension of knowledge production is presented as a counterargument to the view that takes the material world to be the main, if not the sole, justifying factor. Given that the balance tips in favour of the material aspects of experiments, the epistemology of experimentation put forward is here referred to as a 'social' epistemology to highlight the epistemic value of the social dimension of knowledge production in scientific experimentation.

The interfaces between process and product, and between material and social, are taken to provide a complete picture of scientific experimentation that supplies a satisfactory account of how experimenters produce and justify experimental knowledge. The existing accounts of experimentation are then reinterpreted and overlooked problems identified.

SEE'S DESCRIPTIVE AND NORMATIVE FUNCTIONS

The ultimate goal of SEE is to create a framework for the analysis of experimental economics. Given that this framework is built upon existing studies of experimental physics the first doubt that naturally arises concerns the adequacy of a framework built upon a natural science. The arguments that physics might offer experimentation's most successful stories or that economics has indeed attempted to emulate physics are – however favourable – not the arguments advanced here. As mentioned above, the choice of physics is determined by pragmatic considerations. Nonetheless, the study of experimental physics is taken to be relevant for identifying the inner processes and associated epistemic attributes of scientific experimentation. The underlying presupposition is that these inner processes are common to all experimental practice even though how they come about may differ from one experimental field to another. Experimental practice can be generally conceived of as an endeavour that strives to attain a three-way coherence, the epistemic value of which depends upon the levels of materiality and sociality. How the three-way coherences are obtained in actual practice depends upon the problem-situation at hand, which varies both within a single discipline and across disciplines. Thus, the *Social Epistemology of Experiment* is assumed here to be a general framework which can be applied to different units of analysis within different disciplines.

The epistemological framework is simultaneously a descriptive and a normative device for understanding and appraising the practices and products of experiments. The principles underlying the three epistemic arguments of experimentation broadly characterise and justify scientists' practices. However, how close scientists abide by those principles in actual practice is necessarily an empirical matter. This is precisely the object of the four epistemic tests. They are intended to assess the levels of *materiality* and *sociality* of the three-way coherence supporting experimental results.

The epistemological framework enables both the identification and the appraisal of the distinctive features of the experimental method of economics. Its application to experimental economics highlights the specific procedures and arguments used by

economists to produce and establish experimental results. Even though this is not the main concern of this thesis, the distinctive features of experimental practice in physics and economics are made evident. Thus, SEE is not only suitable for the study of particular episodes of experimental practice and resulting products. Rather, it is also adequate for the comparative analysis of different experimental disciplines, sub-disciplines, methods and procedures, as well as experiments.

That the epistemological framework is suitable to both a descriptive and a normative function does not come as a surprise insofar as it is built upon informed accounts of actual scientific practice which capture the fundamental features of the processes in which practising scientists, embedded in scientific communities, produce and establish novel claims to knowledge. This dual capability is a requirement of any normative device insofar as it must take into account the specificity of the particular processes of knowledge production to best evaluate them. This specificity includes the technical requirements of the experimental process of knowledge production (material apparatuses, instruments, procedures,...), the cognitive attributes of individual scientists (perception, memory, computation, impartiality...) and the social characteristics of scientific communities (attitude towards discovery, critical attitude,...).

Finally, it should be mentioned that the account to be presented endorses a fallibilist account of science. The underlying presupposition is that scientists tend to adopt reliable procedures in problem-solving activities that the community of scientists can accept as legitimate and valid. Scientists are concerned about their reputations as competent problem-solvers and the scientific community awards credit to scientists who produce reliable results. However, the history of experiments also shows that experimental results are revised, corrected, and eliminated in the course of practice and along with it their significance to current scientific culture.

1.5 THE EXPERIMENTAL METHOD OF ECONOMICS⁸

The experimental field of economics is mostly an achievement of the second half of the twentieth century.⁹ This explains why economics was considered to be a non-experimental science until very recently and why scepticism remains about the use of the experimental method in economics. Famous quotations by John Stuart Mill

⁸ The expressions 'experimental method of economics' and 'experimental economics' are taken here as equivalents. However, it should be noted that this thesis focuses on the latter, that is, the field of research commonly associated with the label 'experimental economics'. 'Experimental economics' is a well-delimited area of research that can be testified by the 'textbook' definition of the method, the subfields of research and respective sets of experiments and stylised facts (cf. Davis and Holt 1992; Kagel and Roth 1995). However, this characterisation does not exhaust the use of the experimental method in economics. Other uses take place, quite substantially, within the field of Behavioural Economics. Even though the specificity of this set of experiments is not the object of this dissertation, some brief comments are made in Ch. 9.

⁹ There are a few early famous experiments, the most ancient of which is Bernoulli's Saint Petersburg's experiments (1834). But only since the 1950's did experiments begin to be carried out in a systematic fashion. See Roth (1995) for an historical overview.

(1844), Lionel Robbins (1932), Milton Friedman (1953), Richard Lipsey (1979) and Paul Samuelson and William Nordhaus (1985) as to the non-experimental nature of economics are all well-known to economists.¹⁰

Experimental work in economics has increased enormously over the last two decades. A large and growing community of economists conducts experiments in laboratories scattered throughout the world. Hundreds of reports of experimental work are published in top journals such as *Econometrica*, *The American Economic Review* and the *Economic Journal* and in newly specialised journals such as *Experimental Economics*. The discipline has an official association, the *Economic Science Association*. Experimental Economics is a part of the curricula of undergraduate and graduate courses in the leading faculties of economics, and it is a field with one recipient of the Nobel Prize: Vernon L. Smith.¹¹ Thus in general, economics has increasingly become an experimental science.

The methodological and epistemological analysis of experimental economic, however, has lagged far behind. Only in recent years have experimenters begun to reflect upon the experimental method of economics and philosophers of science have begun to select economics as a subject matter of inquiry. Moreover, the methodological and epistemological study of experimental economics is still very much in its infancy. SEE is intended as a contribution to that research program. It does that by supplying the answers to three questions:

- 1) *What are the characteristic features of the experimental method of economics?*
- 2) *How do experimental economists justify their claims to knowledge?*
- 3) *What kind of knowledge does it provide about individual behaviour and economic systems?*

As will be shown, experimenters have been more concerned with persuading the profession about the legitimacy of the experimental method in economics rather than with providing in-depth analysis regarding the role and function of experimental economics. Moreover, the methodological discussion by the experimenters themselves has been characterised by an ill-informed and uncritical use of the available philosophical resources. The implication of this is that most of the epistemological analysis of experimental economics is subject to the problems of the study of experimentation pointed at above. Too much emphasis has been given to theory-related problems and to the end-products of economic experiments, whilst the processes in which the community of experimenters generate their claims to knowledge have been overlooked. It is thus not surprising that the social dimension of knowledge production has not been given sufficient attention. There are, however, a few exceptions to this rule. Timo Tammi (1999a, 1999b, 2003) has

already made an initial incursion into the study of the collective dimension of the production process of experimental economics. Focusing on one of the most debated topics of experimental economics, the preference reversal research program, Tammi has illustrated how the community of economists as a whole can deal with knowledge production problems such as those conveyed by the underdetermination thesis (1999a). However, the neglect of the material aspect of knowledge production (amounting in economics experiments to the participation of human subjects) renders this a fairly partial analysis (as will be demonstrated in Ch. 9).

The special issues devoted to debating the experimental method of economics in the mainstream *Economic Journal* (1999) and in the specialist *Journal of Economic Methodology* (2005) present careful and in-depth reflections on the experimental method of economics, but they are still partial and disconnected contributions to the field. *The Methodology of Experimental Economics* (2005b) by Francesco Guala offers the first thorough philosophically informed and comprehensive analysis of experimental economics. However, the research interest is very much focused on the products of experiments rather than on the processes. In fact, Guala explicitly states that his interest is non-historical, i.e. his concern is with the “sort of knowledge social scientists can collect in the laboratory, regardless of time, place, and context” (ibid, p. 3). The goal is to “put philosophy at use in understanding the rationale of down-to-earth methodological principles”. The overarching theme is inference; inference “from data to phenomena, and from phenomena to their causes within a given experimental setting” and inference “from laboratory circumstances to some real-world situation” (ibid, p. 6).

The *Social Epistemology of Experimental Economics* takes into account the context surrounding the emergence of an experimental field, the methodological arguments, the experimental precepts, debates, and results. Rather than bringing philosophy to clarify methodological and epistemological issues, it applies the knowledge derived from the studies of experimentation to the understanding of experimental economics. The application of the social epistemological framework to experimental economics offers understanding about the experimental process of knowledge production of economics in its technical, cognitive, and social dimensions. In the process, the field’s research programs, conventions, disputes, experiments and results are appraised. However, perhaps the most important contribution of the present research is the presentation of particular aspects of the experimental process of knowledge production (in both its social and technical details), which have been overlooked by existing approaches.

Before presenting the structure of the present research, it should be explicitly recognised that, unlike the studies of the new philosophical approach to experimentation, the study of experimental economics is based on the written reports of experiments, the methodological writings of experimental economists and the methodological and epistemological analysis carried out by the philosophers of science. The direct access to actual experimental practice that is available, for example from participant observation and from conducting interviews with practicing experimenters is therefore very limited. Some direct resources of experimental practice were used nonetheless. These included: the participation, as an experimental subject, in the pedagogical experiments of the course of Experimental Economics of the

¹⁰ See appendix 1.

¹¹ Vernon Smith was awarded the Nobel Prize in 2002 for his pioneering experimental work in economics. In fact, the bulk of this dissertation focuses on Smith’s contributions to the foundation and establishment of experimental economics. The prize was jointly awarded to Daniel Kahneman for the insights he brought from cognitive and experimental psychology into economics (some of which are reviewed in Ch. 9).

University of Amsterdam, in 2001; a visit to the laboratory of the Center for Research in Experimental Economics and Policy Decision-Making (CREED) of the same university and in the same year; the collaboration in a public goods experiment at the research centre Dinâmia, Portugal (Lopes *et al.* 2004). The remaining resources included informal conversations with experimenters and students of experiments, as well as – perhaps most informatively – with sceptical economists.

The method adopted thus contrasts with the new experimentalist methodology the researcher ascribes to. However, this fragility is significantly reduced by the specificity of experimental economics. In fact, the characteristic traits of experimental economics made the present research possible. As this thesis demonstrates, experimental economics is fairly open to the scrutiny of virtually any student of science. This openness is favoured by the publication of:

- 1) substantial information about the actual production processes;
- 2) extensive and in-depth surveys on experimental work undertaken to date;
- 3) experimenters' reflections about the use of the experimental method in economics;
- 4) public debate surrounding the main methodological and epistemological characteristics of experimental economics.

1.5 OVERVIEW

The thesis is organised into two parts. The first part constructs the social epistemology of experiment and the second applies this to experimental economics. The first part is made up of three chapters. Chapter 2 analyses the role of the 'material world' in scientific experimentation while chapter 3 focuses on the epistemic contribution of the social dimension of knowledge production. Chapter 4 presents the four tests proposed for the appraisal of the processes and products of scientific experimentation. The second part of the dissertation is composed of the subsequent five chapters. Chapter 5 provides a brief presentation of the experimental method of economics and analyses one category of experiments – market experiments. Chapter 6 presents and assesses the arguments put forward to justify the use of the experimental method in economics. Chapter 7 focuses on the 'material component' of economic experiments and for that end it studies another category of experiments – game theory experiments. Chapter 8 derives the implications of the 'material' quality of economics experiments. Chapter 9 focuses on the collective aspect of experimental practice in economics, selecting the third grand category of economics experiments – individual behaviour and decision-making experiments. Chapter 10 concludes by applying the *Social Epistemology of Experiment* onto itself.

APPENDIX 1

THE NON-EXPERIMENTAL SCIENCE OF ECONOMICS

[T]here is a property common to almost all the moral sciences, and by which they are distinguished from many of the physical; that is, that it is seldom in our power to make experiments in them (John Stuart Mill [1836] in Guala 2005b, p. 2)

[O]ur belief in [economic generalisations] does not rest on the results of controlled experiments (Lionel Robbins [1932] in Guala 2005b, p. 2)

[U]nfortunately, we can seldom test particular predictions in the social sciences by experiments explicitly designed to eliminate what are judged to be the most important disturbing influences. (Milton Friedman [1953] in Starmer, 1999a, p. 1)

It is rarely, if ever, possible to conduct controlled experiments with the economy. Thus economics must be a non-laboratory science. (Richard Lipsey [1979] in Starmer 1999a, p. 1)

One possible way of figuring out economic laws... is by controlled experiments... Economists [unfortunately]... cannot perform the controlled experiments of chemists or biologists because they cannot easily control other important factors. Like astronomers or meteorologists, they generally must be content largely to observe. (Paul Samuelson and William Nordhaus [1985] in Friedman and Sunder 1994, p. 1)

Economists are sometimes confronted with the charge that their discipline is not a science. Human behavior, it is said, cannot be analysed with the same objectivity as the behavior of atoms and molecules. Value judgements, philosophical preconceptions, and ideological biases must interfere with the attempt to derive conclusions that are independent of the particular economist espousing them. Moreover, there is no laboratory in which economists can test their hypotheses. (*Encyclopaedia Britannica* [1991] in Davis and Holt 1993, p. 4, n. 2)

PART ONE

THE SOCIAL EPISTEMOLOGY OF EXPERIMENT

CHAPTER 2

EXPERIMENT, COHERENCE AND MATERIALITY

2.1 INTRODUCTION

This chapter presents an account of experimental practice that identifies the features intrinsic to the experimental process of knowledge production that help explain and justify experimenters' belief in knowledge experimentally generated.

The account of experiment to be presented is built upon recent studies in the fields of philosophy, history and sociology that can be brought to bear on central methodological and epistemological issues of scientific experimentation. As will become apparent, these are responses to the sociological upsurge in the studies of science that questioned the traditional role of experiment as a neutral scientific adjudicator. This is not surprising, though. It is in times of profound crisis in established scientific worldviews when fundamental methodological and epistemological issues are brought to the fore.

The contributions reviewed are varied, however. Some, more in line with traditional accounts of science, attempt to restore the traditional role of experiment as an objective arbiter, whereas others, inspired by the new sociological approach, attempt to offer an objective account of science while incorporating the social conditioning of the processes of knowledge production. Still, they all conceive experimentation as an important scientific endeavour and experiment as the hallmark of science. Moreover, these studies are well-informed reports of experimental practice based on thoroughly studied historical episodes. However, they constitute but partial contributions to the study of experiment. In this regard, the account of experiment to be presented represents an attempt to offer a comprehensive portrait of scientific experimentation that helps identify and answer methodological and epistemological questions pertaining to the experimental process of knowledge production.

The chapter is structured as follows. Section 2 presents the stylised view of experiment and identifies the methodological and epistemological questions to be addressed. Based on the studies of experiment, Section 3 reconstructs experimental practice as a scientific endeavour that aims at establishing *a three-way coherence* among the main components of the experimental system – the material procedure, the instrumental model and the phenomenal model. It is then argued that the alignment achieved among the components of the experimental system provides scientists with reasons for believing in experimental knowledge – this is the *argument from coherence*. Section 4 illustrates the concept of three-way coherence and section 5 shows how a

three-way coherence is established in actual practice. Section 6 elaborates on the concept of experimental coherence. Because the three-way coherence may carry varied epistemic import, its epistemic significance is scrutinised in section 7. In section 8, the *argument from materiality* asserts that the engagement of the material world in scientific experimentation bestows added epistemic value to the three-way coherence sustaining an experimental result. The three subsequent sections then evaluate the factors that condition the participation of the material world in the experimental process of knowledge production. Section 12 concludes the chapter.

2.2 THE STYLISTED VIEW OF EXPERIMENT

As explained in the introductory chapter, the philosophical study of scientific experimentation is a relatively recent endeavour. The recent interest in experiments emerged as a reaction to the challenges posed by the sociological studies that questioned the privileged status of experimental results. This fact helps explaining, to some extent, the partial nature of the accounts of experimental practice insofar as these have been tailored to answer specific questions. It is therefore not that surprising that there is not, as yet, a standardised definition of a laboratory experiment that the students of scientific experimentation allude to.

No doubt early accounts of experiment already noted the epistemic role of scientists' interaction with nature with the help of instruments (e.g. Bacon) and offered fairly elaborated accounts of the experimental process of knowledge production (e.g. Stuart Mill). But it was not until the last two decades that philosophers of science have put forward systematic accounts of scientific experimentation. Some have highlighted the experimental manipulation of the material world to argue for the reality of theoretical entities (Hacking 1983); others have invoked the notion of 'close systems' (Bhaskar 1975) as a means to acquire knowledge about 'mechanisms' or that of 'nomological machines' (Cartwright 1999) to stress that law-like regularities obtain only (or almost) in the laboratory. Others have emphasised the distinction between the actions performed by the scientists and the interpretation of those actions and respective results to reiterate the potential reproducibility of experimental processes and results (e.g. Radder 1996). This is not to say that there are as many conceptions of experiment as the interests of scientific scrutiny. These different accounts are by no means incompatible. However, a comprehensive view of scientific experimentation is still missing.¹²

More or less implicitly, the studies of scientific experimentation take experimental practice as the scientific endeavour that consists in the production of stable phenomena under controlled circumstances. Four interrelated key concepts are then associated with experimental practice. First, an experiment involves the *production* of a phenomenon for scientific scrutiny. The reason is that interesting phenomena do not generally occur naturally or they do not occur in conditions that

allow for close scrutiny. Second, an experiment involves *control* so that the phenomenon of interest is produced under conditions favourable to scientific scrutiny. This means that in order to producing the phenomenon of interest, experimenters must control for extraneous interferences (the so-called 'noise' or 'background factors') that may have an effect on the phenomenon of interest but which are not the object of inquiry. Third, the control exercised over the experimental conditions aims at *isolation*, that is, the shielding of the produced phenomenon from extraneous influences so that the evidence generated pertains to the object of study and this object only. Finally, the control exercised over the experimental object and the experimental environment accounts for the *stability* of the experimental phenomena. This means that whenever the experimental conditions are reproduced, the phenomenon manifests itself in the same way. In sum, the studies of experiment taken together all support the following broad and encompassing definition:

Laboratory experimental practice consists in the production of stable phenomena under the controlled circumstances of the laboratory to generate evidence relevant to scientific inquiry.

The production of a stable phenomenon in the laboratory is not a trivial accomplishment. It involves a long sequence of operations until the recording of the evidence for the phenomenon can take place. Experimenters need to design, build and implement special experimental devices to produce and explore as yet unknown aspects of the world. In the process unexpected events arise that call for adjustments to the experimental apparatus, refinement of experimental procedures, revision of working assumptions, and so forth. The experimental process of knowledge production ends when the phenomenon of interest is believed to have been successfully produced in the laboratory. In Ian Hacking's words:

To experiment is to create, produce, refine and stabilise phenomena. If phenomena were plentiful in nature... it would be remarkable if experiments didn't work. But phenomena are hard to produce in any stable way. That is why I spoke of creating and not merely discovering phenomena...

There is designing an experiment that might work. There is learning how to make the experiment work. But perhaps the real knack is getting to know when the experiment is working. That is one reason why observation, in the philosophy-of-science usage of the term, plays relatively small role in experimental science. Noting and reporting readings of dials... is nothing. Another kind of observation is what counts: the uncanny ability to pick out what is odd, wrong, instructive or distorted in the antics of one's equipment. The experimenter is not the 'observer' of traditional philosophy, but rather the alert and observant person. Only when one has got the equipment right is one in a position to make and record observations. (1983, p. 230)

From this crude presentation of experimental practice already emerge the fundamental methodological and epistemological questions inherent to scientific

¹² A sample of the portraits of experiment is offered in appendix 2 to this chapter.

experimentation. If experimentation requires the construction of new equipment for the creation and the recording of evidence for unknown phenomena, the central methodological and epistemological questions are, respectively:

- 1) *How do experimenters know they have got the equipment right and/or that they have created the phenomenon of interest rather than an artefact of the experimental procedure?*
- 2) *What are the grounds for belief in knowledge generated by a process in which both the means and the outcomes of that process are at stake?*

Whereas the first question calls for the identification of the set of standards that guide experimenters' actions and decisions in the construction of novel resources for scientific practice, the answer to the second question invites inquiry into the warrant of belief those standards provide. As will become apparent below, the answers to these questions are not as straightforward as one could have expected, not even to account for simple experimental tasks. Here is Hacking's stance while referring to the observation of a cell through a microscope:

Our conviction arises partly from our success at systematically removing aberrations and artifacts... We are convinced about the structures [of the cell] we seem to see because we can interfere with them in quite physical ways, say by microinjecting. We are convinced because instruments using entirely different physical principles lead us to observe pretty much the same structures in the same specimen. We are convinced by our clear understanding of most of the physics used to build the instruments that enable us to see... We are more convinced by the admirable intersections with biochemistry, which confirm that the structures that we discern with the microscope are individuated by distinct chemical properties too. We are convinced not by a high powered deductive theory about the cell – there is none – but because of a large number of interlocking low level generalisations that enable us to control and create phenomena in the microscope. In short, we learn to move around in the microscopic world. (1983, p. 208-9)

Hacking suggests that confidence in experimental results is provided by a variety of scientific resources and practices which together provide reasons for believing that the phenomenon created in the laboratory is the phenomenon of interest rather than an artefact of the experimental procedure. Experimental knowledge relies on the possibility of physically interfering with the object under scrutiny, on available instruments grounded in different theories, on well-established theories of phenomena, on low level generalisations, and so forth. These resources and practices are the basis of two arguments Hacking put forward to justify the validity of experimental results.

As we have seen in chapter 1, the 'argument from intervention' asserts that the confirmation of expectations while interfering with the object of scrutiny strengthens belief in both the proper operation of the experimental apparatus and the observations made with it (e.g. when looking at a cell through a microscope, the observation of a change in the colour of the cell after injecting fluid strengthens the

belief in the proper operation of the microscope and in the observations made with it, *ibid.*, p. 189-90). The 'no-miracle argument' states that the generation of the same results with different apparatuses provides reasons for believing those results because different apparatuses are associated with different 'backgrounds' which render very implausible the explanation that the systematic 'errors' of the different apparatuses generate exactly the same results. Conversely, the attribution of the same evidence to a correct operation of various instruments is fairly understandable and reasonable. In the case of observing through the microscope, "it would be a preposterous coincidence if, time and again, two completely different physical processes produced identical visual configurations which were, however, artifacts of the physical processes rather than real structures in the cell" (*ibid.*, p. 201).

However satisfactory these arguments may seem, they do not really offer a comprehensive view of experimental practice that enable answering the central methodological and epistemological questions of experiment. If not for other reasons, experimenters may not have at hand different alternative apparatuses nor reasonable expectations about what a proper intervention amounts to. In the next section it is argued that these arguments, resources and practices, individually and jointly, establish relations of coherence and that coherence is the underlying principle that guides experimental practice and grounds belief in experimentally generated knowledge.

2.3 EXPERIMENTATION AS THE FORGING OF A THREE-WAY COHERENCE

I take experimental practice as a scientific endeavour that aims at establishing a 'three-way coherence' among the three main components of the experimental system – the 'material procedure', the 'instrumental model' and the 'phenomenal model'. This conception and related terminology is borrowed from the sociologist Andrew Pickering (1989, 1995b) whose account of scientific experimentation most directly bears on the two guiding questions of the present research. Nonetheless, as will be shown below, the conception of scientific experimentation as the forging of relations of coherence and the presupposition that coherence is the underlying principle supporting experimental knowledge are also pervasive, more or less explicitly, in other studies of experiment.

To better illustrate what is entailed by the expression 'three-way coherence' and how it is actually attained in practice, some terminological clarification is in order. In the remainder of the manuscript expressions such as 'problem-situation', 'problem-solving', 'problem-solution', 'scientific culture' and 'resources of practice' are recurrently used. The expression 'problem-situation' refers to the problem the experiment is conceived and implemented to solve and respective background knowledge, that is, the underlying scientific culture that provides meaning to the problem. On this view, problems arise when scientists' expectations are disappointed in the light of what is taken as established beliefs. The expression 'problem-solving' refers to the experimenters' actions undertaken to find a solution to that problem and the expression 'problem-solution' the solution arrived at. Finally, 'scientific

culture' encompasses the established items of science or the 'resources of practice' which include instruments, techniques, theories, meta-level theoretical presuppositions, metaphysical commitments, and so forth. The scientific culture is established by scientific communities and respective institutions, such as universities, laboratories, scientific associations and publishing systems, according to shared standards of what constitutes good scientific practice (the role the institutions of science play in experimental practice is addressed in section 2.7 below and it is the central topic of chapter 3).

THE THREE-WAY COHERENCE

Pickering describes experimental practice as a scientific endeavour that consists in the manipulation of the three components that make up an experimental system – the 'material procedure', the 'instrumental model' and the 'phenomenal model' – until a 'three-way coherence' obtains (1989, 1995b). According to Pickering, the material procedure comprises "experimental action in the material world: setting up the apparatus, running and monitoring it in the laboratory". The instrumental model refers to "the experimenter's conceptual understanding of how the apparatus functions" and the phenomenal model refers to "the conceptual understanding of whatever aspect of the phenomenal world is under investigation" (1989, p. 276-77).¹³

Figure 2.1 below presents a schematic representation of the practice leading to the three-way coherence and the experimental result it supports. Experimental practice gets started with the identification of a 'problem-situation'. This often results from a clash between established knowledge and recent results of scientific practice that conflict with it. 'Problem-solving' then proceeds by finding a feasible way of investigating the material world. This requires searching through the available items of scientific culture the material and the conceptual resources that may be mobilised and adapted to that end. When experimenters believe that they have finally identified a potentially workable material apparatus whose operation can provide the information needed about the material world, a two-way coherence is achieved between the instrumental and the phenomenal model (cf. figure 2.1a).

The experimental practice then continues with the material implementation of the instrumental model. Because experimenters cannot fully anticipate the outcome of their actions, experimental practice gives rise to surprising results that call for further intervention in the material world in order to gain understanding about the

unexpected results. Material and conceptual manoeuvrings then feed back upon the various components of the experimental system (cf. figure 2.1b). In the process, new material apparatuses and techniques may be tried out and the understanding of how the object of scrutiny is to be investigated may be substantially revised. When the implementation of the instrumental model finally produces results that can be interpreted by the phenomenal model, a 'three-way coherence' obtains (cf. figure 2.1c above). When the three-way coherence obtains the experimental process of knowledge production ends. It provides experimenters sufficient reasons to put an end to their practice. As Pickering puts it:

The output phenomenal model, and the fact it carries, must be right, it seems, because it is implied by the material procedure and the instrumental model which lie behind it. Of course, as a plastic resource, there is no guarantee that this particular instrumental model is right, but it fits so nicely between the fact and the material procedure that it is hard to doubt. On the other hand, there is no guarantee that the material procedure is the correct one, material procedures are plastic, too, but... and so on. (1989, p. 280)¹⁴

It should be noted that the coherent material procedure, instrumental model and phenomenal model are substantially modified versions of their originals. They are the outcome of material and conceptual manipulations carried out during the experimental process of knowledge production. (These transformations are captured in figure 2.1 by the change of the shapes of each of the components of the experimental system during the various stages of the process of knowledge production. The replacement of the dashed lines by the straight lines, in turn, depicts the stabilisation of the coherent components)

When the process of knowledge production ends the components of the experimental system crystallise in their final versions. The instrumental model then provides the description of how the material apparatus is to be implemented and operated to generate information about the aspect of the material world under investigation. The instrumental model, by conveying the conceptual understanding of the material procedure, endows the experimental result significance as a reliable outcome, that is, as a result that was adequately generated. The phenomenal model interprets the results produced by the material procedure and thereby endows the outcome of the experiment significance as a carrier of knowledge about the material world. The material procedure, in turn, acquires significance as an adequate means of investigating that particular aspect of the material world. However, the 'three-way coherence' has only achieved a temporary stabilisation. It may be destabilised in subsequent practice.


¹³ How these concepts translate into the components of the experimental systems of economics is postponed until chapter 5. For now it suffices mentioning that the 'material procedure' in economics refers to the procedures used to create economic behaviour and socioeconomic institutions in the laboratory. The 'instrumental model' refers to the reasoning that justifies the effectiveness of experimental procedures in producing that behaviour and those institutions. The 'phenomenal model' refers to the interpretation of the observed behaviour and of the institutions. The 'material world' of economic experiments therefore consists in the object of scrutiny of economics, i.e. the behaviour of the experimental subjects and the socioeconomic context wherein that behaviour takes place.

¹⁴ To Pickering any product of science is plastic in the sense that it is amenable to manipulation. However, as he acknowledges, the plasticity of scientific resources is limited in the sense that they afford a limited number of manipulations. For instance, conceptual models are constrained by the boundary conditions imposed by the respective theoretical frameworks and an instrument is constrained by what is taken as an admissible operation. The plasticity of the resources of practice is addressed in detail in section 2.9.


FIG. 2.1. The Experimental Process of Knowledge Production



a) Establishing how a given phenomenon is to be investigated by experimental means: the two-way conceptual coherence between the instrumental model and the phenomenal model.



b) The forging of the three-way coherence by material and conceptual manoeuvring.



c) The end of the experimental process of knowledge production: the attainment of the three-way experimental coherence.

The phenomenal model contains a characteristic that is absent in the other two experimental components. The phenomenal model is both the input and the output of experimental practice. The experiment is designed and implemented to produce a phenomenon of interest of which experimenters have only a faint idea. This idea is the original input phenomenal model. But insofar as the goal of the experiment is to attain knowledge about that phenomenon, the input model is by necessity incomplete, inaccurate or even incorrect. If experimenters had full knowledge about the phenomenon there was no purpose in conducting the experiment in the first place. Thus, the original phenomenal model is to be transformed in the course of experimental practice until experimenters are confident that they have made sense of the phenomenon under investigation. This transformation accompanies the material and conceptual manoeuvring of the experimental system until the material procedure generates a phenomenal model that adequately accounts for the experimental outcomes. To put it briefly, whereas the input phenomenal model conveys an expectation about the aspect of the world to be investigated, the output

phenomenal model carries the fact that has been determined in the course of experimental practice. It is a fact that was not known, at least not in its full extension, prior to practice. This fact is the ultimate experimental result that once established gains autonomy from the material and conceptual work that generate it. It becomes an independent fact (cf. picture 2.1c above).

Following Pickering's conception of experimental practice, I can now provide a first sketch of the answers to the methodological and epistemological questions of experiment. The *Social Epistemology of Experiment* takes coherence to be the central principle guiding scientific practice and the grounding for belief in the results of experiment. Coherence is a three-way alignment between the components of the experimental system, the material procedure, the instrumental model, and the phenomenal model. This alignment forms a unified, organised and structured experimental system such that experimenters believe they have understood it.

The three-way coherence grounds belief in the experimental results by virtue of the alignment achieved between the components of the experimental system. By fitting so nicely with one another, each component finds support and supports the other component parts of the system and thereby the result conveyed by the output phenomenal model. Thus, when scientists achieve a three-way coherence they are convinced that they have produced the phenomenon of interest or that they have controlled the influence of extraneous factors so that the phenomenon produced is no artefact of the experimental apparatus. However, this is a result that can only be determined in subsequent practice. The output phenomenal model has yet to be subjected to the scrutiny of the scientific community (see Ch. 3). In sum, the social epistemology of experiment yields the *argument from coherence* that asserts that the three-way coherence achieved between the component parts of the experimental system provides scientists with reasons for believing in the experimental results.

COHERENCE IN THE STUDIES OF EXPERIMENT

The conception of scientific experimentation as the forging of relations of coherence among the components of the experimental systems is supported by other studies of experiment, notwithstanding the differences in jargon.

The sociologist David Gooding thinks of experimentation as a learning process that “may show convergence of successive material arrangements (the apparatus) and successive construals (or models) of manipulation of and with apparatus, and its outcomes... to a stable state in which all three are mutually compatible”. The significance of this stable state is that it “indicates that the process is locally convergent, that is, producing sense or order that can be reproduced”. Otherwise, “it fails to make experience intelligible and reproducible” (1990, p. 166-7; see also 1992).

In the philosophical camp, Hacking came also to conceive experimentation as a practice that achieves the ‘mutual adjustment’ of the ‘matériel’, ‘ideas’ and ‘marks’ until they are brought into a relation of ‘consilience’ which accounts for the ‘self-vindication’ of experimental results. The matériel, ideas and marks correspond, roughly, to Pickering's material procedure, instrumental model and phenomenal model (cf. Hacking, 1992, p. 32). Also the philosopher Hans Radder conceives the experimental

process as involving “the material realization and the theoretical description or interpretation of a number of manipulations of, and their consequences for, the object and the apparatus, which have been brought into mutual interaction” (1995, p. 58, emphasis omitted). The view that the coherence established among the multiple and the heterogeneous items of experimental culture gives scientists confidence that they have produced and understood the laboratory phenomenon is thus well supported by various accounts of experiment (see appendix 2 for more examples). To better illustrate what is entailed by the expression ‘three-way coherence’ and how it is actually attained in practice, Pickering’s account of *The Hunting of the Quark* is summarised in the next section (1981; 1984; 1989; 1995b, Ch. 3).

2.4 THE HUNTING OF THE QUARK 1: ESTABLISHING THE THREE-WAY COHERENCE¹⁵

In the early 1960s it was established knowledge that the most elementary particles of matter were ‘hadrons’ carrying electric charges of integral multiples of e (the charge of the electron). In 1964, Murray Gell-Mann and George Zweig challenged this view when they conjectured that instead of being elementary particles, hadrons were composites of more fundamental entities, the ‘quarks’. In addition, they attributed to quarks the property that they carry third-integral electric charges of either $e/3$ or $2e/3$. This theoretical conjecture organised the recent findings produced by the new generation of particle accelerators that had identified over two hundred hadrons. Even though this was a mere theoretical conjecture, a group of physicists at the University of Genoa, led by Giacomo Morpurgo, welcomed the Gell-Mann and Zweig’s suggestion and adopted the search for fractionally charged objects as their *problem-situation*.

The search for quarks began in conceptual practice. The group reasoned that since, as it was believed, all matter carried electric charge in integral multiples of e , fractionally charged quarks, should they exist, would be easily recognised. The classical electrostatic theory and the oil drop experiment of the American physicist Robert Millikan were the main items of *scientific culture* the Genoa group resorted to. From the laws of classical electrostatics they knew that the force (F) on any macroscopic body placed in a uniform electric field (E) is the product of the charge (q) on the body and the magnitude of the electric field ($F = q * E$). From Millikan’s experimental work they knew that the electric charge (q) on a body could be deduced from measurements of the force (F) on the body suspended in a known electric field (E).

Given that fractionally charged entities should be rare objects were they to exist, the team reasoned that much larger samples of matter than those used by Millikan should be inspected. This created the first difficulty the group encountered. Millikan’s

original experiment used an electric levitation system which was deemed inadequate to sustain larger samples of matter. This meant that a new detector had to be constructed for that purpose. Nonetheless, they still believed that the principles of Millikan’s experiment could be applied to the new device. They then decided to build a new material apparatus that used a dielectric suspension system instead of the Millikan’s electric system. The conceptual possibility of measuring charges on large samples of matter by applying an electric field to particles suspended in dielectric liquid was therefore the team’s first *instrumental model*. This instrumental model achieved a two-way coherence between the instrumental and the phenomenal model in that the former showed a feasible way to investigate whether the material world would be one of integral or third integral charged particles.

As it could be expected, the first attempt at implementing the detector with the dielectric system produced behaviour that the experimenters failed to make sense of. The electric charges obtained were neither integral nor third integral. The failure encountered represented a frustration to a prior expectation formed in conceptual practice. The experimenters were willing to accept one among two phenomenal models: either integral or third-integral charged particles. The mismatch between the material procedure and the phenomenal model was attributed to the material apparatus. Given that the major difficulty of the Millikan-type apparatuses consists in creating an isolated homogeneous electric field, so that the measurements obtained can be attributed only to the charge on the body, the group interpreted the results as being caused by spurious effects produced by an exchange of charges between the sample and the suspending liquid. Consequently, the revision of the conceptual understanding of the material procedure solved the inconsistency. The liquid suspension system was deemed inadequate for the purpose at hand. A more appropriate system, they conjectured, would have to create a vacuum to avoid the interaction between the sample and the suspending system.


The experimenters decided to design and implement a new detector: the magnetic levitation electrometer. The new detector used a magnetic suspension system to make the grains of matter levitate in a vacuum. Electric forces would be applied to the magnetic field (E) and the force (F) on the grains measured. Having the values of E and F the electric charges (q) on the matter would then be derived by calculation. This time, the material implementation of the new apparatus generated integral charges. As a result, the *material procedure* interpreted through the new *instrumental model* produced results consistent with one of the *phenomenal models* the team was willing to accept. The *three-way coherence* between the material and the conceptual items of practice yielded the result that there are no fractionally charged particles – this was the *output phenomenal model* of the quark experiment.

Of course, this is an overly simplified account of the practice of the Genoa team. The group encountered other difficulties in their search for quarks. For instance, in the course of practice a grain was identified which manifested a strange behaviour. The grain displayed the same movement in the same direction both when the electric field was applied and when it was reversed. This could indicate the presence of an unknown entity that was causing the change in the sign of the electric charge. But the perplexing result was again attributed to inhomogeneities in the electric field, in particular, to interferences

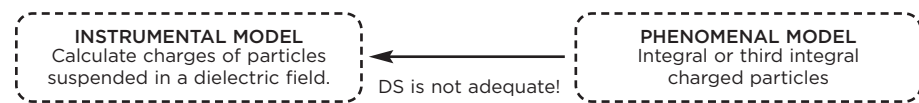
¹⁵ The account given here is taken to be a faithful description of the experimental program on elementary-particle physics which does not commit the ongoing analysis to Pickering’s conclusions. This will become apparent as Pickering’s case study is used to reveal the flaws of his main arguments (in this and in chapters 3 and 4). It should be noted that Pickering’s critics also acknowledge the account’s accuracy. For instance, Gingras and Schweber state that “Pickering’s book [1984] is a first-rate history and explanation of the subject, accurate and readable at the same time” (in Hacking 1999, p. 65).

of electric field gradients on the grain. Under this reading, the results became intelligible and the three-way coherence was restored. The material realisation of the experiment interpreted through the instrumental model yielded the result that the charge of any body is an integral multiple of e . In 1966, G. Gallinaro and G. Morpurgo published in the *Physics Letters* the result that they had been unable to find quarks. Below is the schematic representation of the experimental process that generated that result.


FIG. 2.2 The Hunting of the Quark



a) The two-way conceptual coherence.




Dielectric system is not adequate!



MLE in working fine!

b) The forging of the three-way coherence.



c) The attainment of the three-way experimental coherence.

The hunting of the quark did not end with the successful completion of this experiment. The Genoa group continued to search for quarks. The group wanted to search in more massive grains which required the design of new material procedures. As they experimented on ever more massive samples of matter, the same pattern of practice followed: continual deformations of material procedures and conceptual models so as to eliminate the inconsistencies encountered in practice and thereby establish a three-way coherence. Whenever the three-way coherence was attained, scientists were confident that sufficient support had been gathered for their practices and the results were subsequently submitted to public scrutiny. In 1977, G. Gallinaro, M. Marinelli and G. Morpurgo reported in *Physics Review Letters* that they continued to fail finding evidence for quarks. In 1981 Morpurgo decided to end the research program.

2.5 ESTABLISHING RELATIONS OF COHERENCE

The conception of experimental practice as a scientific endeavour that consists in the design, implementation and manipulation of a material apparatus and in the ongoing revision of related conceptual models until a three-way coherence obtains is a crude simplification. The three-way alignment involves a variety of resources and practices to establish the overall fit between the material procedure, the instrumental model and the phenomenal model. In this section I show that this is especially the case when new resources of practice must be constructed from scratch to solve the problem at hand. This so because the construction of novel resources requires the establishment of a more intricate network of coherent results in their support. Conversely, experiments that resort to well-established resources do not need to undertake such work because the reliability of the resources has already been attained in prior practice. That is, previous practice has already established these resources as adequate solutions for relevant problems. From this it follows that *the three-way coherence supporting an experimental claim to knowledge involves a more or less complex web of relations of coherence depending on the reliability of the resources of practice.*

To illustrate how a three-way coherence is build up in practice and how this constructive work depends on the reliability of the resources of practice I now have recourse to Allan Franklin's 'epistemological strategies' of experiment. These are here taken as means of generating partial and local level relations of coherence that together help constructing the three-way coherence.

Franklin presents the strategies of experiment as "arguments designed to establish, or to help establish, the validity of an experimental result or observation" (1989, p. 437; see also 1986, 1990a). But he notes that the list provided is not "exclusive or exhaustive" and that no single argument, or subset of arguments for that matter, is meant as "necessary or sufficient condition for rational belief" (1989, p. 459).

Like Hacking's arguments (section 2.2), Franklin's list of strategies is not explicit about what is actually accomplished by the strategies that provides the grounds for rational belief, nor what their relative epistemic import is.¹⁶ In this section I argue that each item in the list may be considered as a *coherence procedure*, that is, a means of establishing a coherent relation between resources of practice and thereby the

alignment between material procedure, instrumental model and phenomenal model. On this view, there is no a priori need to set the number or the combination of procedures required for the attainment of the three-way of coherence. What matters is that that relation be established and this may be achieved by any possible combination of procedures. The more coherent the relations that are established, the higher the epistemic value of the three-way coherence.

To better illustrate how the coherent relations are established in actual practice, the epistemological strategies of experiment are reviewed in some detail (cf. Franklin 1989). The strategies of experiment are:

1. *Independent confirmation by different apparatuses*: scientists are justified in believing their observations when different apparatuses produce the same observational patterns.¹⁷
2. *Indirect validation of a single apparatus*: in cases where an observation can only be made by using one kind of apparatus, scientists are justified in believing in that observation if the same apparatus reproduce other phenomena observed by different physical processes.
3. *Observing the predicted effect of intervention*: observing the predicted effect strengthens belief in both the proper operation of the experimental apparatus and in the observations made with it.¹⁸
4. *Prediction of the absence of phenomenon*: one has justified belief if an effect disappears when it is predicted that it will.
5. *Using an apparatus based on a well-corroborated theory*: evidence supporting the theory of the apparatus grounds belief in the apparatus and thence in the observations made with it.
6. *Using properties of phenomena as validations*: justified belief is grounded on the properties of the observed phenomena themselves that could not be plausible explained by a malfunction of the apparatus.
7. *Well-corroborated theory of phenomena*: the validity of observations can be strengthened if they can be explained by a well-accepted theory.

¹⁶ A similar criticism is made by Giora Hon (2003, p. 180-1). He argues that Franklin's list lacks an "overall guiding principle" and therefore "do not form a comprehensive view of experiment" which may provide a "coherent generalization of experiment" that reflects "the inner processes, the epistemic dynamic processes, inherent in experiment that may help answering methodological and epistemic questions" (2003, p. 181). Hon proposes the elimination of error as the overall guiding principle of experiment. The main argument of this chapter is that "the overall guiding principle" of experiment is the establishment of a three-way coherence which provides a "comprehensive view of experiment" that answers to fundamental methodological and epistemological questions. As will become apparent below, error elimination is here treated as one among various coherence procedures.

¹⁷ This corresponds to Hacking's 'no-miracle argument' alluded to in section 2.2.

¹⁸ This is Franklin's version of Hacking's 'argument from intervention' referred to in section 2.2.

8. *Elimination of plausible sources of error and alternative explanations of the result*: the demonstration that the observations cannot be caused by other plausible factors or interpreted by alternative explanations increases belief in the reasons provided for the observations.¹⁹

9. *Calibration and experimental checks*: the calibration of experimental apparatuses by reference to known phenomena argues for the proper operation with the apparatus and in favour of the results obtained.

10. *Statistical validation*: statistical arguments provide additional reasons for believing in experimental results.²⁰

The strategies of experiment just reviewed are very illustrative in showing that the highest degree of confidence occurs when the experimental result can be grounded on alternative and well-established material apparatuses. This is the case when distinct material procedures, each with a well-defined instrumental model describing the appropriate operation with the material apparatus, yield the same phenomenal model (strategy 1).²¹ The reproduced result is then accepted with the force of certainty because it is a highly replicable and robust result.²² Reproducing results that are also generated by other apparatuses reinforces the reliability of the

¹⁹ This strategy invokes a recurrent theme in the accounts of experiment. On this view, confidence in experimental results is built by the progressive elimination of 'error', 'background effects', or 'noise', until scientists are confident that the observations refer only to the phenomenon of interest. A representative view of this strategy is Galison (1987, 1997) that is analysed below. A more comprehensive view that extends beyond the elimination of spurious effects is provided by Hon (1989, 2003). Deborah Mayo (1996) provides an elaborate 'error-elimination' epistemology of experiment with an emphasis on the use of statistic arguments. As mentioned in the introduction, the 'error-elimination' accounts of experiment endorse a negative epistemology that hides the actual achievements of experimental practice which the social epistemology of experiment tries to bring to the fore.

²⁰ This strategy has grown in importance with complexity. The use of large-scale and expensive material apparatuses, such as those characteristic of high-energy physics, renders the construction of new devices or their adaptation for new problem-situations unfeasible. Under these circumstances, much of the experimental work leading to an experimental conclusion relies on work done after operating with a relatively stable apparatus. It relies instead on data-processing procedures that search through the mass of experimental data the fraction that is to constitute evidence for the phenomenon of interest. Data analysis therefore attempts to deal with the backgrounds that could not be controlled or eliminated by material means (c.f. Galison 1987, p. 263-6).

²¹ Franklin notes that well-established instruments allow 'direct observations' from nature. Quoting Dudley Shaper: "X is directly observed if (1) information is received by an appropriated receptor and (2) that information is transmitted directly, i.e. without interference, to the receptor from the entity X (which is the source of the information)" (in Franklin 1986, p. 165). Simon Shaffer uses the term "transparency" to describe the attribute of an instrument when it is treated as a reliable transmitter of 'nature's messages' (1989). Well-established items of scientific practice have also been labelled 'black boxes' to convey the autonomy of scientific results from their context of production. The present research takes a well-established instrument to be a material apparatus that has achieved a two-way coherence with an instrumental model that describes its proper operation for specific purposes. (This topic is addressed in detail in section 4.3)

²² The notions of replicability and robustness are the standard epistemic criteria of experimental results. The notion of replicability invokes the stability of a result produced with the same experimental system and the notion of robustness invokes the stability of the result in the face of variations introduced to the original experimental system. To put it in another way, whereas a well-established three-way coherence sustains a replicated result, various three-way coherences sustain a robust result. These notions are analysed in more detail in section 4.4.

material apparatus to applications beyond the cases where they are substitutes (strategy 2).

The use of an experimental apparatus that has not yet achieved the status of a reliable instrument demands a tighter relation between material procedure, instrumental model and phenomenal model. The definition of the phenomenal model, or models, must then be prior to the design and implementation of the experiment. The supporting evidence for one of those models will provide confidence in the operation with the apparatus and the predicted result (strategies 3 and 4). The weaker epistemic value of this three-way coherence may be enhanced by additional coherent relations with the theory of the apparatus (strategy 5), the observations obtained (strategy 6), or the theory of the phenomenon under investigation (strategy 7). The epistemic worth of the experimental result can be further improved if other plausible and alternative interpretations are also accounted for (strategy 8).

When the available knowledge does not yet allow for the creation of the phenomenon of interest, or its production is unfeasible, a better-known apparatus may be used to generate a surrogate for it. However, this requires careful scrutiny so as to make the experimental result to bear on the target phenomenon (strategy 9). This is an epistemically weak result until technical progress finally allows the generation of the phenomenon of interest. In any case, statistical analysis can be carried out and further assurance provided that the experimental data are significant to the problem-situation at hand (strategy 10).

The coherent procedures described above show that the web of partially coherent relations established in current practice depends on the availability of well-established resources of practice. In the case of well-established material apparatuses, it is enough for scientists to be assured that the apparatus has worked properly. The understanding of what an appropriate operation is has already been established in prior practice. Under these circumstances, a *two-way coherence* between the material procedure and the respective instrumental model suffices to warrant belief in the experimental result. When feeble apparatuses are used, a tighter set of partially coherent relations is required between the material apparatus and the theory of the apparatus, or between the experimental evidence and the theory of the phenomenon, and so forth. At any rate, a complex array of coherences must sustain each and every experimental result. The question is whether these coherences had already been established in prior practice. If not, they still have to be obtained.

On this view, then, well-established items of scientific culture presuppose previously made coherences, coherences first established in their production and coherences made in subsequent practice as the items of culture became resources to solve other scientific problems. Conversely, new resources have still to be established by forging new relations of coherence. At any rate, *an epistemically strong result is a result that presupposes an intricate web of relations of coherence among various items of scientific culture regardless of whether these relations were forged in the past or in present practice.*

In this section I showed how the epistemic value of the three-way coherence depends on the epistemic value of the items of scientific culture. In section 2.7 I will show that the epistemic value of the three-way coherence also depends on the

participation of the material world in knowledge production. Before that the concept of three-way coherence will be further refined.

2.6 THE CONCEPT OF EXPERIMENTAL COHERENCE

Coherence, understood in a broad sense as compatibility between held beliefs, is a widely epistemic esteemed value in science. Coherence is an attribute that should characterise single theories, the relation between theories and high-ordered principles, the relation between theories and evidence and, ultimately, all established knowledge. It is therefore not surprising that philosophers have adopted the notion of coherence in their theories of justification.

Coherence theories take justification of any belief to be dependent on inferential relations to other beliefs and eventually to the overall system of beliefs held by the believer.²³ Different concepts of coherence have been proposed according to the *elements of coherence*, i.e. the objects that cohere, and the *relations of coherence*, i.e. the coherent connections between those objects. In the coherence theories of justification the elements of coherence are propositions and the relations of coherence invoke logical, inferential or explanatory relations. Relations of coherence may therefore come in different guises and degrees. The strongest notion of coherence incorporates relations of nomological necessity or logical entailment. Other relations of coherence include, for instance, probabilistic relations of various kinds (e.g. the likelihood of a belief in relation to others), explanatory relations, and relations of consistency.

Explanatory coherentism has become popular in recent years and it has been developed to address issues of theoretical assessment. Paul Thagard (1989, 1992), for instance, developed a theory of explanatory coherence for the appraisal of competing hypotheses. On this view, the explanatory hypothesis that should be accepted is the one that best coheres with other beliefs. This means that the hypothesis must have superior explanatory breath (in Whewellian sense as expressed by the notion of 'consilience') and must organise the largest amount of observations, while avoiding contradiction and competition with established empirical and theoretical propositions. Thagard demonstrated the historical accuracy and the psychological plausibility of his theory by showing that it accords with well-known historical episodes of theoretical change (e.g. the acceptance of Lavoisier's oxygen theory, Darwin's evolution by natural selection, wave theory of light, etc.). In these accounts, however, coherence pertains to relations between theoretical and empirical propositions taking the latter as given.

The coherence account of scientific experimentation put forward here concerns the production of empirical propositions and the making of relations of coherence between the components of the experimental system: the material procedure, the instrumental model and the phenomenal model. Rather than establishing relations

²³ Coherence theories of justification are a twentieth century philosophical achievement. The seminal contributions are due to F. H. Bradley, Bernard Bosanquet and Brand Blanshard. More recent positions are advocated by Gilbert Harman, Nicholas Rescher, Keith Lehrer and Laurence Bonjour. See Bonjour (1985) for a survey of philosophical ideas about coherence.

between ready-made propositions, the relations and the propositions are both part of the same production process whereby scientists try to make sense of their actions and resulting outcomes. The relations of coherence therefore concern the interpretation of scientists' actions and of the results of these actions.

The making of the three-way coherence is the result of a complex activity that extends beyond the establishment of inferential relations between sets of propositions. The contributions of the cognitive sciences are in this respect more fruitful to shed light on the processes whereby scientists try to make sense of their new interventions in the material world. Cognitive scientists have already highlighted the importance and ubiquity of *mental representations* in the formation of belief. These not only mobilise representation tools such as concepts, schemas, images and analogies, but also *mental procedures* such as deduction, induction, matching, mapping and adaptation (cf. Thagard 1996). All these mechanisms are called for as scientists try to make sense of their active interaction with their objects of scrutiny and try their tentative interpretations against the material world and those of other scientists.

Moreover, the production of the experimental three-way coherence is an *interactive process* whereby the components of the experimental system are continuously revised until a coherent resolution emerges. In this sense, the experimental notion of coherence resembles the notion of 'reflective equilibrium' as applied in ethics (cf. Rawls 1971). In ethics, the notion of 'reflective equilibrium' has been used as a criterion on the basis of which moral judgements about particular issues are to be assessed. This examination consists in evaluating the extent to which these moral judgements cohere with other beliefs about similar cases as well as other beliefs about moral and factual issues. As a result of this evaluation, the moral judgements are reflected upon and revised until they cohere with other beliefs. A reflective equilibrium is reached when the inquirer has no inclination to revise the elements of the belief system any further because they have a high degree of coherence. Seeking coherence among the widest set of beliefs therefore contributes to the revision and adjustment of beliefs. In the process, prior beliefs may be modified and new beliefs added to the system of beliefs. On this view, a belief is justified if it coheres with the rest of beliefs on due reflection and after appropriate revisions have been operated throughout the system of beliefs. A person who holds a belief in reflective equilibrium with other relevant beliefs is then justified in believing in that principle or judgment.

Finally, the making of relations of experimental coherence is part of a social process. On the one hand, the making of the three-way coherence is a social endeavour undertaken by collectives of scientists whose activity takes place in well-organised and structured scientific institutions. On the other hand, the three-way coherence is a crucial element for persuading the wider community of scientists that the experimental process of knowledge production is reliable and therefore that there are objective grounds for belief in the experimental results.²⁴ The analysis of the social

dimension of the experimental process of knowledge production and of its epistemic role is, however, postponed until the next chapter.

I can now synthesise the notion of the three-way coherence of the social epistemology of experiment and its epistemic role in the experimental process of knowledge production.

The experimental three-way coherence:

- *It is a three-way alignment between the material procedure, the instrumental model and the phenomenal model.*
- *When it obtains the material procedure interpreted through the instrumental model produces facts within the framework of the phenomenal model.*
- *It is worked-out during the process of knowledge production through a wide range of conceptual and material practices.*
- *The practices that produce it bring about a plurality of mental representations and processes whereby scientists make sense of the world and of their intervention in it.*
- *It is a collective endeavour that strives for the acceptance of the experimental results by the larger scientific community.*

The three-way coherence provides scientists with reasons for believing in the experimental results, because:

- *The alignment achieved between the components of the experimental system supports each single component and the system they make up and thereby the experimental result conveyed by the output phenomenal model.*

When the three-way coherence obtains the process of knowledge production ends, because:

- *Scientists believe that they have produced the phenomenon of interest; or they believe that they have controlled the influence of extraneous factors so that there is no need to introduce further amendments to the experimental system.*
- *Scientists believe that they are justified in believing in the results supported by the three-way coherence and that the scientific community will believe in them, too.*
- *However, confidence on experimental results may be destabilised in subsequent practice.*

So far I have argued that from the accounts of experiment coherence emerges as the unifying principle that guides experimental practice and grounds belief in experimental results. But simply alluding to the mutual support of the components of the experimental system does not suffice to justify belief in experimental results. Firstly, as it was already noted, the epistemic value of the three-way coherence is variable. It depends on the reliability of the items of scientific culture resorted to, that is, the extent to which these have been established in prior practice. Secondly,

²⁴ Mäki (1995) provides a coherence account of rhetorical persuasion along similar lines. On this account, the intensification of a persuadee's belief in a given proposition depends on the proposition's plausibility which depends on the relation of coherence this proposition establishes with the set of premises from which it is derived and that is believed by the persuadee.

the three-way coherence must support a significant claim to knowledge. It must carry relevant information about the aspect of the material world under investigation. The implication of this is that further arguments must be provided to justify the confidence afforded by the three-way experimental coherence.

2.7 THE EPISTEMIC VALUE OF COHERENCE

It is pretty consensual in the studies of experiment that the epistemic value of experimental results is conferred to by the participation of the material world in knowledge production. Indeed, this participation is the distinguishing feature of scientific experimentation as compared to other modes of scientific inquiry. The participation of the material world in knowledge production is epistemically relevant because the actual properties of the aspect of the world under scrutiny can play an active role in the production of knowledge about them. This role is manifested by the potential of the experimental process of knowledge production for generating unanticipated results.

Pickering's conception of scientific practice as a *dialectic of resistance and accommodation* (1990b, 1994a, 1995b) offers an illustrative account of the way by which the material world participates in knowledge production. According to Pickering, experimental practice is 'dialectic' because it consists of scientists' actions in the field of material agency and in the reactive responses of the material world to those actions. The dialectic is made of 'resistance and accommodation' because, on the one hand, the material world 'resists' scientists' expectations and, on the other hand, scientists attempt to 'accommodate' this resistance in subsequent practice. The material world poses obstacles to scientists' expectations because the reaction of the material world to human actions cannot be fully anticipated. This is, as noted above, where the interest of experimenting resides, i.e. to learn a not fully understood aspect of the material world. Thus, in the course of human intervention in the material world unanticipated problems arise that have to be addressed if scientists want to find solutions to their scientific problems. This means that in experimental practice the resistance of the material world represents a failure in fulfilling an expectation (*viz.* the anticipated problem-solution) and scientists' accommodation to material resistances represents an attempt to meet an expectation, albeit a revised one (*viz.* the problem-solution arrived at).

The 'dialectic of resistance and accommodation' is meant to stress the unpredictability of science, since practice brings about an unpredictable sequence of actions, and the intricacy of human and material agencies, insofar as experimental practice is determined by both human and nonhuman agencies each of which is generative of the other. That is, scientists' actions are responses to material resistances and the resistances encountered result from human actions. From this it follows that the results of experimental research cannot be known in advance. Regardless of scientists' expectations, the engagement with the material world raises unanticipated problems which lead to a succession of unforeseen actions whose results are unpredictable.

Pickering's account is targeted against the radical reading of the sociological approach to science that attributes a prominent role to the social aspect of knowledge

production (Pickering 1992, 1994a, 1995b), namely the role of 'interests' (cf. Barnes 1977, 1982; Shapin 1982; Mackenzie 1981) or that of 'social constraints' (cf. Bloor 1983). Pickering's point is that scientists' purposes, goals and plans are not immune to the 'dialectic' for they can as well be revised when confronted with material resistances. No doubt Pickering misrepresents the implications of the strong programme of the sociology of scientific knowledge (c.f. Bloor 1991), but the relevant point, for the present purpose, is that the results of experiment involve the participation of the material world and this participation has epistemic value. It promotes the revision of scientists' beliefs. In Pickering's words:

[I]nitial plans do not determine outcomes. It [the dialectic] emphasizes instead the redirection of modelling vectors in the course of dialectics of resistance and accommodation taking place in fields of material and disciplinary agency, dialectics that are themselves constitutive of the detailed products of practice. Further, these dialects can entail the reformulation of plans that, anyway, collapse in machinic captures of material agency and in the achievement of interactive stabilisations of multiple material and conceptual cultural elements. The products of scientific practice, then, including knowledge (as well as machines, instruments, disciplines, social relations, and so on) are significantly detached by the mangle²⁵ from the subjectivity of human agents... this detachment from the intentional structure of human agency through encounters with material and disciplinary agency²⁶ (themselves proper to no individual subject) is the basic sense of the objectivity of science that the mangle makes available to us. (1995b, p. 195)

On Pickering's view, not only are the products of science objective because they cannot be anticipated, they are also objective because "the successful captures, framings, and interactive stabilisations that characterise the objective contents and products of science are hard to come by; their achievement is difficult and uncertain" (p. 195-6). In sum, the participation of the 'material world' is the major factor contributing to the epistemic value of the results of experiment. It generates new phenomena and thereby contributes to the growth of knowledge. The epistemic value of the participation of the material world in knowledge production is now presented in the *argument from materiality*.

The engagement of the material world in knowledge production conveys epistemic value to the three-way coherence and the experimental result it supports. The three-way coherence is epistemically significant because it is the outcome of a process that produces new phenomena when the material world resists to scientists' expectations. Insofar as the three-way coherence accommodates the resistances


²⁵ The mangle is "a convenient and suggestive shorthand for the dialectic" that "conjures up the image of the unpredictable transformations worked upon whatever gets fed into the old-fashioned device of the same name used to squeeze the water out of the washing". In a broader it sense it refers to "the worldview... that sees science... as an evolving field of human and material agencies reciprocally engaged in the play of resistance and accommodation" (1994a, p. 567).

²⁶ Disciplinary agency is a procedural item of scientific culture, in particular, "the sedimented, socially sustained routines of human agency that accompany conceptual structures" and it plays "an analogous role in conceptual practice to that of material agency in material practice" (1995b, p. 29).

of the material world, it assures scientists that they have succeeded in making sense of the new phenomena. The experimental three-way coherence is a carrier of a novel claim to knowledge.

In figure 2.3 below the added epistemic value of the three-way coherence conveyed by the participation of the material world in knowledge production is depicted by the reduced dimension of the elliptic representing the phenomenal model. The smaller shape of the phenomenal model represents the more limited range of admissible solutions that results from the participation of the material world in knowledge production.

FIG. 2.3 The Material World and the three-way experimental coherence



2.8 THE EPISTEMIC VALUE OF THE PARTICIPATION OF THE MATERIAL WORLD

The justification of the epistemic value of the three-way coherence needs to take also into account the social dimension of knowledge production, which is the subject matter of the next chapter. Before that, the epistemic role of the material world in knowledge production needs closer inspection. We still need to investigate the extent to which the material world is given a chance to participate in the process of knowledge production and the relevance of this participation to the growth of knowledge. This is the natural question to raise insofar as the potential for generating novel knowledge depends on the actual participation of the material world.

To do that, I will first look at two extreme situations to help illuminating the most common cases. The first situation pertains to experimenting with rigid experimental systems while the second concerns experimenting with plastic experimental systems. In the first situation, the experimental system does not provide substantial room for scientists' manoeuvring. The process of knowledge production is highly dependent on the scientific culture that guides the selection of the scientific problems and provides the means for solving them. Under these circumstances, problem-solutions are fairly anticipated by the available resources of practice insofar as their domains of applications and the kinds of results they yield are to some extent already known prior to practice. In the second situation, the higher degree of manipulability of the experimental systems renders problem-solutions more vulnerable to the influence of scientists' actions.

Scientists may either revise the material procedure so as to make it conform to favoured conceptual models or they may by further conceptual manipulation explain the resistances away. This means that plastic experimental systems constrain the participation of the material world by allowing for a high degree of human intervention. Rigid experimental systems, while they do not easily allow accommodating material resistances, they do narrow down the range of action of material agency. In either way the epistemic role of material agency is or may be reduced.

What follows from this is that plastic experimental systems tend to generate problem-solutions sustained by weakly established three-way coherences. Rigid experimental systems generate results well-supported by three-way coherences, but the phenomenal models they generate do not have a high potential to carry new knowledge. These are two extreme cases in which the three-way coherence is either epistemically weak or epistemically redundant.²⁷ They, however, reveal the importance of analysing the factors that may have an effect on the epistemic value of the three-way coherences sustaining experimental results.²⁸

THE RELIABLE BUT REDUNDANT THREE-WAY COHERENCES

Reliable but redundant three-way coherences are more likely to characterise problem-solutions of 'normal science' (Kuhn, 1970[1962]). 'Normal science' is the scientific practice exercised within the borderlines defined by the 'paradigm' that supplies the criteria for selecting problem-situations and the tools to solve them. This dependence constrains the problem-situations to research questions that can be solved with the available resources and assessed by the standards of the paradigm. Scientists are therefore compelled to conform to these standards if they want their results to be accepted by the scientific community. That is, in 'normal science' scientific culture provides the means of knowledge production and the standards by reference to which new knowledge is evaluated. But it constrains novelty by generating results that are by and large already embedded in extant scientific culture. As a result, the discovery potential of problem-solving is seriously limited. This is the 'essential tension' inherent to scientific practice that Kuhn (1977) identified.

The limited potential for discovery of 'normal science' is well depicted by Kuhn. The problems that cannot be expected to be solved or that are not solved by the paradigm's resources are not conceived of, or they are ignored or rejected as external to the purview of the paradigm, and consequently rejected as adequate scientific inquiries.²⁹ This is how Kuhn framed this limitation:

²⁷ Redundancy is here used in a broad sense to convey the requirement that science should add new items to established knowledge and extant scientific culture. The worry is that scientific practice becomes the mere exercise of instrumental virtuosity rather than an actual problem-solving activity.

²⁸ That coherence may possess different levels of epistemic value has been noted in other coherence accounts. Hare (1973) argued that the coherent resolutions obtained through the process of 'reflective equilibrium' may be vulnerable to individuals' intuitions and therefore ineffective in promoting the revision of beliefs.

²⁹ Different concepts of a paradigm have been found in Kuhn's *Structure of Scientific Revolutions* (Masterman 1970). The conception that is followed here is the conception of the paradigm as an 'exemplar' for scientific practice (Kuhn 1970, p. 186).

One of the things a scientific community acquires with a paradigm is a criterion for choosing problems that, while the paradigm is taken for granted, can be assumed to have solutions. To a great extent these are the only problems that the community will admit as scientific or encourage its members to undertake. Other problems, including many that had previously been standard, are rejected as metaphysical, as the concern of another discipline, or sometimes as just too problematic to be worth the time. (1970, p. 37)

The problem-solutions that conflict with the paradigm are deemed erroneous and a failure of the scientist rather than a fact of nature, consequently, they are set aside and neglected. The implication of this is that normal scientific practice consists to a great extent in “an attempt to force nature into the preformed and relatively inflexible box the paradigm supplies” (op. cit., p. 24). Such a practice not only does not promote novelty but it also leaves it unnoticed when it does occur.

No part of the aim of normal science is to call forth new sorts of phenomena; indeed those that will not fit the box are often not seen at all. Nor do scientists normally aim to invent new theories, and they are often intolerant of those invented by others. Instead, normal-scientific research is directed to the articulation of those phenomena and theories that the paradigm already supplies. (ibid)

Kuhn does not perceive this limitation as particularly troublesome. The accumulation of failures in problem-solving eventually generates new problem-situations that ultimately give rise to a new basis for the practice of science, i.e. a new paradigm. Science shifts from paradigm to paradigm. A paradigm's potential is fully explored and a new paradigm emerges to replace the old one. In the meantime, normal science extends the knowledge of those facts that the paradigm displays as interesting.

The view endorsed by the present research is less optimistic, however. On the one hand, dominant scientific cultures do have an active role in determining the direction of science and thus the development of new ‘paradigms’. On the other hand, the scientific cultures wherein scientists are educated create cognitive limitations to the exploration of new domains of research. But for now this issue is left aside. In the second part of this thesis illustrations will be provided that show how the cognitive and psychological predispositions of scientists prevents scientists from perceiving and accepting novelty when facing it and how scientific communities and institutions create resistance to change.

THE UNRELIABLE THREE-WAY COHERENCES

Whereas ‘normal’ scientific practice constrains the choice of problem-situations, ‘extraordinary’ science raises questions about the reliability of problem-solutions. The reliability of the problem-solutions of ‘extraordinary’ science is likely to be small due to the absence of well-established resources and standards for practice. While in ‘normal’ science social, cognitive and psychological factors reinforce well-established standards (e.g. preference for solutions that accord with dominant cultures, resistance to change, deference to authority, sensitiveness to peer pressure), in ‘extraordinary’

science these factors tend to create rupture with well-established beliefs (e.g. preference for solutions disagreeing with dominant cultures, preference for change, radicalism).

In extraordinary science, scientific endeavours are more sensitive to the influence of scientists’ prior beliefs regarding the problem at hand and how a solution to that problem is to be obtained. In the absence of shared standards of appraisal it becomes difficult to decide what a legitimate test for a given problem-solution might be. Under these circumstances, scientists may prematurely put an end to their research projects whenever they succeed in obtaining the anticipated solution while overlooking or neglecting conflicting results. As a result, the problem-solutions arrived at are more significantly the realisation of the scientists’ preferred solutions to the selected problem-situations.

In sum, whereas ‘normal’ problem-situations limit the production of novelty while producing reliable results, ‘extraordinary’ problem-situations are conducive to novelty but the scientific results it generates are not as reliable in that they are not generated by socially validated means of knowledge production.

The two types of problem-situations portray extreme categories of research. In actual scientific practice the difficulties they allude to are not that clear-cut. On the one hand, scientists’ priors depend on scientific culture, which is the basis of belief formation, and scientists are irremediably dependent on the available resources of practice which provide the means of problem-solving. On the other hand, for any scientific culture there is a degree of freedom concerning the choice of the sorts of problems to embrace and the kinds of solutions to favour. Thus, scientists’ actions are by and large guided by attempts to realise envisaged solutions. Nonetheless, the degree of scientists’ freedom increases as problem-situations deviate from extant culture.

To conclude, *the potential of the experimental systems for generating novel knowledge is lower the more the problem-situations conforms to established scientific culture. But the more they do so, they also produce more reliable results. Conversely, autonomous problem-situations potentially produce novel but yet to establish results.* Thus, in laboratory experimentation *there is a trade-off between the influence of scientific culture and scientists’ priors or that between the production of socially validated results and those with a high potential for novelty.*

THE EXPERIMENTAL PROBLEM-SITUATIONS

The present work focuses on problem-situations that lie in between the two extremes considered, for in science there is no such thing as starting from scratch and scientists’ minds are no *tabula rasa*. However, the epistemic problems of interest here approximate those of extraordinary science. The most interesting epistemic questions of scientific experimentation pertain to those cases where scientists produce the material means of knowledge production anew (albeit from extant resources) to explore new domains of research. When experimenting in novel domains, knowledge production is inevitably guided by speculative ideas about how to design, modify and adapt the material apparatus as well as about the phenomenon the experimental system is to produce in the laboratory. Under these circumstances, the agency of the scientists in knowledge production is high. Consequently, a clash between the

expectation and the unanticipated result may be more easily resolved by attempting to obtain support to prior beliefs than by revising them. This may be achieved by further manipulation of the material apparatus so as to materially eliminate the conflict, or by explaining the effect away by conceptual manoeuvring.

This vulnerability of the experimental problem-solutions was clear in the description of the coherence procedures listed above. When experimenters can resort to well-established resources of practice, their behaviour is predominantly guided by routine (e.g. the two-way coherence associated with the operation of a reliable material apparatus, cf. strategy 1). When the means are produced along with the results of practice scientists' expectations and actions assume a more significant role in knowledge production (e.g. confirmation of predictions, cf. strategy 3). This is not to deny the importance of experimental work that revolves around the manipulation of well-established resources (e.g. for improving the precision of measurements, refinement of instruments, checks on prior results, etc.). The point is simply that these problem-situations have a weaker potential for generating novel knowledge. I now illustrate how the participation of the material world may be constrained when experimenting with plastic and with rigid experimental systems.

2.9 THE HUNTING OF THE QUARK 2: THE PLASTIC EXPERIMENTAL SYSTEMS

According to Pickering the resources of practice have a fairly limited role in constraining the behaviour of scientists. He depicts experimental activity as a practice that consists in the design, implementation and manipulation of plastic experimental systems. On his view, the plasticity of the resources of practice together with the unpredictable character of the agency of the material world has the potential to generate an indefinite number of valid solutions. That is, the plastic experimental systems produce or have the capability to produce various solutions none of which may be singled out as superior to the others. To better illustrate this point an account of another episode of the quark search is in order.³⁰

The search for the quark was not exclusively pursued by the Genoa team. Another set of experiments was undertaken under the supervision of William M. Fairbank, at Stanford University, California, which generated evidence for quarks. Similarly to the Genoa's team, the Stanford group sought to search for quarks in an improved Millikan-type experiment. In 1965, they constructed an apparatus that used a magnetic levitation electrometer (analogous to the one used by the Genoa experiment) and an oscillating electric field (also adopted by the Genoa team in subsequent experiments). Based on Fairbank's experience on superconducting technology, the group also used superconducting quantum interference detectors to measure the charge of superconducting niobium spheres floating in a magnetic field. In their first set of measurements they obtained values consistent with the existence

of quarks, though they attributed the results to spurious charge forces caused by the apparatus. However, subsequent attempts to identify and eliminate spurious effects were unsuccessful. The evidence of quarks remained. After twelve years of research, the Fairbank group finally reported that they had found evidence of quarks. In 1977, in the *Physical Review Letters*, they reported six niobium spheres with zero charge, a fractionally charged sphere and a last one that carried a fractional charge which was lost on remeasurement in a subsequent run. Regardless of the care taken in twelve years of research, the experimental results were subject to severe criticism. In response, they introduced new modifications to their material apparatus to account for alleged sources of spurious effects. The improved apparatus continued to generate evidence for fractional charges as reported in the same review in 1979.

Given that the Stanford results were in conflict with extant results (those of the Genoa and other research groups) they were subject to harsh criticism. In a public debate between Fairbank and Morpurgo, published in 1977 in *Physics Today*, Morpurgo identified possible sources of spurious effects to which Fairbank provided counter-arguments. Their disagreement revolved around issues such as the measurement scale (maximum amplitude vs. rates of change of amplitude), the flexibility of the capacitor plates (movable plates vs. movable objects), the elimination of the patch effects by using copper plates; and the importance of having spin samples (not possible in the Stanford experiment). More corrosive, however, was Morpurgo's attack on the credibility of the Stanford group by expressing surprise about the latter's exact findings – they obtained third integral charges while no evidence had been found for some other fraction of the electric charge. This criticism insinuated that the Californian data had been calibrated by reference to the phenomenal model that postulated the existence of isolated quarks. Fairbank replied by resorting to a statistical argument: "Errors were determined... by a computer-derived least squares fit using all of the data and not merely the sample of the data" (in Pickering 1981, p. 228).

This is not to imply that the Stanford experiments were more plastic than the Genoa experiments. The Morpurgo team, too, seemed to have calibrated its results by reference to the opposite model. As described above (section 2.4), whenever the team encountered inconsistent values, subsequent practice either materially corrected those values or explained them away. Nor does it suggest that the practices of these experimenters were guided by deceit. The plasticity of the experimental systems simply rendered the results vulnerable to the effect of the phenomenal models they were designed to explore. Furthermore, both teams invoked the same coherence procedures and attained a three-way coherence in support of their respective phenomenal models. Thus, two experimental processes that had recourse to the same or similar items of scientific culture generated conflicting results – quarks in one case, no quarks in the other.

The Hunting case study is used by Pickering to show that for any problem-situation there might be indefinite and conflicting problem-solutions and that there is no "neutral" way to adjudicate among alternative solutions. In the first place, there were no distinctly different features in the practices of the two groups. Both attempted to produce credible results that could face up the criticism of others.

³⁰ The same argument for the parity-violation experiments can be found in Pickering (1990a).

And in the case of the Stanford group, a significant amount of subsequent revisions were introduced in response to those criticisms. Secondly, the charges levelled against the Stanford experiments concerned practical details of the experiment rather than major violations of shared standards of good experimental practice. The point I want to make, however, is that these experiments demonstrate the extent to which highly plastic resources of practice nullify the agency of the material world. In the quark research program, different teams of researchers generated conflicting results and the practice of each team continued to produce results consistent with the phenomenal models they first obtained. Commitments as to the (non)existence of isolated fractionally charged particles remained unchanged.

2.10 THE RIGID EXPERIMENTAL SYSTEMS

Peter Galison is more in line with Kuhn's 'normal science' and therefore his account of experimental practice can be more easily depicted as experimenting with rigid experimental systems. Therefore, Galison's contribution provides the foil against which Pickering's view might be confronted.³¹ Indeed, to Galison the rigidity of the resources of practice is what conveys epistemic value to experimental results.

Galison's account constitutes a response to the sociological studies of experiment that see the absence of an algorithm that indicates the termination point of the experimental process a symptom of the irrationality of experimental enterprise. Galison's response is that the items of scientific culture impose 'constraints' on experimental practice and by so doing aid scientists in narrowing down the "alternatives of what the experimentalist takes to be reasonable beliefs and actions" until the decision to end an experimental project is reached (1987, p. 246). That is, the 'constraints' of scientific culture define and delimit the field of inquiry and thereby contribute to the closure of the experimental process of knowledge production. By performing the role of an experimental algorithm, the constraints also limit the effect of extrinsic values onto the experimental process of knowledge production.

Galison distinguishes between three categories of items of scientific culture which differ in their relative capability to work as a 'constraint' of scientific practice.³² To put it briefly, *long-term constraints* comprise metaphysical commitments, meta-level theoretical and experimental presuppositions that single out particular phenomenal domains and adequate modes of inquiry, instrumentation and techniques. These constraints are widely accepted and assumed by the scientists in the design and

interpretation of experiments in such a way that "if a result contradicted it, the experimentalist would look again at his instruments and procedures" (op. cit., p. 246). These constraints are long lasting because they survive the constitution and conclusion of shorter-term projects. *Middle-term constraints* consist of theoretical and experimental 'programmable goals' that are "attached to specific institutions and people" (op. cit., p. 249). They provide the conditions for a particular experiment by setting the analysis within the framework of a particular theory and the material practice within the realm of a specific device taken as valid for the investigation at hand. According to Galison, middle-term constraints reduce the number of alternatives taken as reasonable grounds for belief by directing scientists towards the most relevant background factors and the most adequate means to tackle them. Finally, the application of particular theoretical models or phenomenological laws and specific experimental practices introduces *short-term constraints* on experimental practice. These constraints are short-lived insofar as "several different theories or models may be compatible with the broad constraints of any given program" and therefore any of them may be discarded and replaced by others. But they are just as valuable heuristics as the other constraints are. Whereas long-term constraints help defining the problem-situation and middle-term constraints build confidence in the means of problem-solving, short-term constraints help deciding when the solution has been arrived at. According to Galison short-term constraints do that by the formulation of precise predictions. The experiment stops when experimenters find empirical support for their predictions. It gives them confidence that they have managed to produce the phenomenon of interest. Otherwise, they go on trying to figure out what went wrong in their experiments (op. cit., p. 252).

Galison's account is very illuminating in showing the dependence of experimental knowledge production on scientific culture and on scientists' anticipations. Scientific culture provides criteria that help select interesting problem-situations, the proper resources to solve them, the parts of the experimental system to be revised in the face of 'background factors' and by so doing it delimits the range of possibilities that appear reasonable until closure is finally attained. Scientists' expectations, on this account, operate as a short-term kind of constraint in the formulation of particular predictions that determine the end of the experiment whenever expectations are met. Even though only the positive contributions of scientific culture and scientists' priors are highlighted, Galison reveals some awareness of their detrimental effect:

All of these factors cast experimentalists into a continual dilemma. On the one hand, without a theory, they frequently will have neither a guiding qualitative sense of where the interesting physics lies nor any quantitative prediction about how large the effect ought to be. With their goal shrouded in darkness they may not be able to find the effect sought or to dissociate it from disturbances. In this sense theories – better, the many layers of theory – provide essential, self-consciously imposed constraints on acceptable data... On the other hand, given a quantitative prediction, the experimentalist is eventually forced to declare (at least implicitly) that there are no more systematic errors. *This stopping place is, naturally enough, all too dangerously often the predicted result.* (op. cit., p. 73-4, emphasis mine)

³¹ Indeed, Galison (1987) is to some extent a response to Pickering (1984), in particular, to the latter's failure in addressing the role of tradition in controlling scientists' practices. For a more in-depth analysis see the debate between the two authors in Buchwald (ed., 1995).

³² Given the centrality of the constraints of practice in his analysis it is surprising that Galison does not provide a theory of constraints that justifies their relative epistemic import. According to the account provided here rigid resources of practice are well-established resources whose epistemic value lies on a dense web of coherences forged in past practice (cf. section 2.5).

However, Galison does not take the detrimental effect of the ‘constraints’ very seriously. He argues that “it emphatically does not follow that expectations are always met” for the resources of practice “while restrictive are not absolutely rigid” (op. cit., p. 257). That is, they are not rigid to the extent “that makes it impossible (or unreasonable) for the physicist to start with one set of beliefs and come... to experimental conclusions contradicting the starting assumptions” (op. cit., p. 258). Implicit in his argument is also the agency of the material world, or the ‘stubbornness’, as he puts it, of the effect despite all efforts to eliminate it (p. 259). Thus, according to Galison the epistemic value of experimental results relies both on scientific culture, which counteracts experimenters’ subjective commitments, and the agency of the material world that cannot be fully controlled by neither of them.³³

Galison’s account suggests, similarly to Pickering, that the ‘stubbornness’ of the material world can attenuate the inevitable dependence of knowledge production on both established knowledge and scientists’ expectations. But Galison adds the ‘resistance’ of scientific culture that counteracts the exacerbated interference of scientists’ predispositions. However, as Galison eloquently depicts, the ‘agency’ of scientific culture may restrict the role of material agency. Problem-situations deeply embedded in existing scientific culture restrict the participation of material agency by limiting scientific problem-solving to well-defined domains of research whose problem-solutions are by and large anticipated by established knowledge.

2.11 PLASTIC RESOURCES AND RIGID EXPERIMENTAL SYSTEMS

Hacking offers an intermediate position. Like Galison he thinks that the resources of practice possess varying degrees of rigidity.³⁴ Among the listed resources of laboratory practice, ‘topical hypotheses’ are among the most revisable elements of the system. They are propositions about the phenomenon of interest that are derived from the more stable ‘systematic theory’ about the subject matter and the experimental results. Material apparatuses also vary in their degree of rigidity, ranging from the off-the-shelf apparatuses to devices constructed anew for specific purposes. According to Hacking, procedures of ‘data assessment’ or ‘reduction’ are not called into question by the experimenters, they are taken for granted. ‘Data analysis’, in turn, depends on the research ‘question’. The most rigid resources of practice are the ‘styles of scientific reasoning’ which are “the expectations about what the world is like and practices of reasoning about it... [that] govern our theories and

our interpretation of data alike”.³⁵ These elements are so rigid that they do not even figure in the list of resources of practice because scientists do not use them. Consequently, they can be ignored because “experimenters do not change ideal conceptions of the universe in the course of, or at any rate because of, experimental work” (1992, p. 44-54). Finally, Hacking also accepts the revisability of any of the resources of practice – “[though] things are “established” before the experiment... none of these is established in the sense of being immutable” (op. cit., p. 50).

Thus, according to Hacking the degree of plasticity of the resources of practice is much more slender than Pickering admits. In contrast to the latter, who considers the possibility of having an indefinite number of solutions to the same problem-situation, Hacking says that “it is extraordinarily difficult to make one coherent account, and it is perhaps beyond our powers to make several” (op. cit., p. 55). The point is that regardless of the rigidity or plasticity of scientific culture, the associations established are hard to come by so that they hardly allow scientists to produce results that fit their prejudices. To put it in another way, even though the resources of practice are plastic, the coherences they bring about are not, nor are the results they yield. In fact, the coherence achieved is not only sufficient to support the experimental results, but it also long-lived.

The stability of experimental knowledge derives from the exclusive reinforcement of the various components of the experimental system. The material apparatus, the phenomenon and the theory that account for it evolve in a mutually self-enforcing way that is not destabilised by work done in other fields. Once a coherence relation is established, theory and instruments remain valid in their data domain. According to Hacking, the stability of experimental knowledge is, in fact, the characteristic trait of ‘laboratory science’. It underlies the ‘self-vindication’ of the results of laboratory sciences.³⁶ This means that Hacking takes scientific experimentation as a far less contingent endeavour than Pickering does. What is contingent is the scientific culture (which he dubs the ‘*form*’ of a branch of scientific knowledge) that arises out of a historical process, but the scientific results (which he dubs the ‘*content*’ of science) are determinate (1999). Scientific culture evolves in time and determines the classes of problem-situations considered as relevant, but once a problem-situation is defined, the problem-solution is determined. Scientific culture in Hacking’s account is therefore the most deterministic. Not only does scientific culture constrain the production of knowledge, but it also generates unique and stable solutions to the problem-situations. New cultures arise, new problem-situations emerge, and old problem-solutions remain valid in their domains of application.

Hacking does not perceive as problematic the constraining effect of scientific culture to the generation of relevant problem-situations or that of scientists’

³³ Galison also argues that the epistemic value of experimental results can be further improved by “increasing the directness of measurement and the growing stability of the results” (1987, p. 259, emphasis in original). While the increase in directness points to the successful attainment of new relations of coherence with the more active participation of the material world (e.g. by materially eliminating a previously estimated error), the increase in stability points to the reproduction of results with different experimental systems.

³⁴ But neither does he supply an explanation for the differentiated levels of rigidity.

³⁵ The laboratory sciences’ style is characterised by two main metaphysical presuppositions: “(a) the expectation that we find out about the world by interfering with it... (b) the expectation that nature “herself” works that way, with forces and triggering mechanisms and the like, and in general a master-slave mode of interaction among her parts” (1992, p. 50).

³⁶ Hacking’s analysis is not applicable to all experimental sciences, though. It applies to ‘laboratory sciences’ “whose claims to truth answer primarily to work done in the laboratory” (1992, p. 33).

presuppositions to the production of reliable problem-solutions. Hacking's analysis, at best, applies to the finished episodes of science, especially to those in which the theories that account for the phenomena have already been developed and stabilised in their respective domains. The SEE instead attempts to provide a framework of analysis that accounts for the various stages of the knowledge production process, including those for which alternative interpretations of the laboratory phenomena have still to be scrutinised and discriminated.

2.12 CONCLUSION

In this chapter scientific experimentation was presented as a scientific endeavour that aims at establishing a three-way coherence among the main components of the experimental system – the material procedure, the instrumental model and the phenomenal model. Two arguments were presented to justify the epistemic value of experimental knowledge. The *argument from coherence* asserts that the alignment of the components of the experimental system provides scientists with reasons for believing in experimental results. The *argument from materiality* was then introduced to reinforce the epistemic value of coherence. It asserts that the engagement of the material world in knowledge production endows epistemic value to experimental knowledge. The resistance of the material world to established knowledge and to scientists' prior beliefs renders the three-way coherences non-trivial achievements. But it was noted that the participation of the material world in knowledge production might still be constrained by scientific culture and scientists' presuppositions. The argument from materiality ultimately turned out insufficient to account for the epistemic value of coherence.

The studies of experiment reviewed showed that. However, none took it very seriously. They all assume that scientific culture can be superseded and scientists' priors frustrated. An epistemology of experiment with normative aspiration must take a different stance, however. Granting the capability of scientific culture and that of material agency to offset the negative effect of prior beliefs in knowledge production, and that of material agency to avoid the restraining effect of scientific culture, the mere invocation of these capabilities does not suffice to justify experimental knowledge. A further argument must be provided to that end. This task is undertaken in the next chapter. It will be shown that the social dimension of knowledge production reinforces the relevance of both the coherence and the materiality arguments of experiment. The *argument from sociality* is thereby introduced.

APPENDIX 2

THE VARIOUS PORTRAITS OF EXPERIMENT

EXPERIMENTS AS SUBSTITUTE SYSTEMS

Such experimental systems [laboratory] having been first designed and then built, are systems that are not, at the end, examined for their own sake. They are artificially designed and constructed substitute systems, controlled mini-worlds that are directly examined in order to indirectly generate information about uncontrolled maxi-world outside the laboratory – such as economic systems and behaviour 'in the wild' (they can also be used for generating information about theories about such systems and behaviours)... this means such experimental systems are representatives of some real non-experimental systems: they are material models of aspects of the rest of the world. This is not to say experimental set-ups are unreal other than in the sense of being artificial, man-made. They are bits-in-the-world and at the same time bits-about-the-world...

In order to construct a substitute system, the experimenter modifies and purifies some relevant chunks of the world. Experimental models isolate fragments of the world by way of controlling the rest of it – or more accurately: for other potentially causally significant things – by causally manipulating those other things, thereby preventing them from interfering with what causal forces operate and what happens in those substitute systems. (Uskali Mäki 2005, p. 306-7)

EXPERIMENT AS A TECHNOLOGY

[D]oing experiments is more an activity to produce *technical effects*, which can be described appropriately as engineering rather than as a scientific activity, properly speaking, as a construction of machines rather than as an inquiry into nature, as an attempt to produce artificial processes or states rather than as a search for true sentences. [Peter Janich 1978, p. 11 emphasis in original]

EXPERIMENT AS NOMOLOGICAL MACHINE

[L]aws of nature obtain – to the extent that they do obtain – ... on account of the repeated operation of a system of components with stable capacities in particularly fortunate circumstances. Sometimes the arrangement of the components and the

setting are appropriate for a law to occur naturally, as in the planetary system; more often they are engineered by us, as in a laboratory experiment. But in that case, it takes what I call a *nomological machine* to get a law of nature...

What is a nomological machine? It is a fixed (enough) arrangement of components, or factors, with stable (enough) capacities that in the right sort of stable (enough) environment will, with repeated operation, give rise to the kind of regular behaviour that we represent in our scientific laws. (Nancy Cartwright 1999, p. 49-50 emphasis added)

EXPERIMENT AS A CLOSE SYSTEM

The experimental scientist must perform two essential functions in experiment. First, he must trigger the mechanism under study to ensure that is active; and secondly, he must prevent any interference with the operation of the mechanism. These activities could be designated 'experimental production' and 'experimental control'. The former is necessary to ensure the satisfaction of the antecedent (or stimulus) conditions, the latter to ensure the realization of the consequent, i.e. that a closure has been obtained....

Only if the mechanism is active and the system in which it operates is closed can scientists in general record a unique relationship between the antecedent and consequent of a law-like statement. The aim of an experiment is to get a single mechanism going in isolation and record its effects. Outside a close system these will normally be affected by the operation of other mechanisms, either of the same or of different kinds, too, so that no unique relationship between the variables or precise description of the mode of operation of the system will be possible. In general, experimental activity requires a degree of plasticity of the antecedent (stimulus) and circumambient conditions to human manipulation and control. (Roy Bhaskar 1975, p. 53).

THE COMPLEXITY OF EXPERIMENTAL PRACTICE

Experiments may be portrayed as involving manipulations of elements in the material world under conditions of control....

Even in those "simple" cases of experiment, theoretical beliefs about the way the world works, personal knowledge of how the experiment can be made to work, and networks of trust that indulge credibility of results and subsequent inferences to parallel situations all figure essential elements of the experimental experience. The ways of modern science are often even further from this vision, for they rely on a huge technological interface between scientist and nature, both in making the intervention and in capturing and assessing outcomes of the intervention. This interface consists of tools to carry out the intervention, instruments to detect the resulting changes in phenomena and yet more instruments to clarify, assess, and present the evidence in a form that can be read by humans. (Mary Morgan 2003, p. 216)

EXPERIMENT, CONTROL AND MEASUREMENT IN ECONOMICS

The fundamental objective behind a laboratory experiment in economics is to create a manageable "microeconomic environment in the laboratory where adequate control can be maintained and accurate measurement of relevant variables guaranteed" [reference omitted]. "Control" and "measurement" are always matters of degree, but there can be no doubt that control and measurement can be and are much more precise in the laboratory experiment than in the field experiment or in a body of Department of Commerce data. (Vernon Smith 1982, p. 930)

CHAPTER 3

THE ARGUMENT FROM SOCIALITY

3.1 INTRODUCTION

In chapter 2 experimental practice was reconstructed as a scientific endeavour that aims at establishing a three-way coherence among the main components of the experimental system: the material procedure, the instrumental model and the phenomenal model. The *argument from coherence* asserts that the three-way coherence achieved between the component parts of the experimental system provides scientists with reasons for believing in experimental results. The mutual reinforcement of the components of the experimental system assures scientists that they have produced and made sense of the phenomenon of interest. But it was also noted that the simple attainment of a three-way coherence does not provide sufficient justification for the experimental result supported by it.

The *argument from materiality* was then introduced to justify the epistemic value of the three-way coherence. It asserts that the engagement of the material world in knowledge production has a potential to generate new claims to knowledge because it frustrates expectations, from the viewpoint of both established knowledge and scientists' prior beliefs. However, the argument from materiality does not yet satisfactorily account for experimenters' confidence in experimental results. The problem is that the participation of the material world may be severely constrained. Rigid experimental systems limit the agency of the material world to fairly confined problem-situations and plastic experimental systems limit the agency of the material world by allowing a high degree of scientists' intervention in the process of knowledge production.

This chapter introduces the *argument from sociality*. This argument together with the argument from coherence and the argument from materiality helps providing a more complete account of how experimenters produce and establish knowledge. The chapter is organised as follows. Section 2 presents the argument from sociality which is illustrated with another episode of the *Hunting of the Quark*, in section 3. The Social Epistemology of Experiment is then compared with other social epistemologies in section 4. Section 5 presents the *argument for criticism* that attempts to dissolve a potential tension that the SEE brings to the fore, the need of a common background for the exercise of effective criticism. Section 6 rounds up the arguments of the SEE. Section 7 concludes.

3.2 THE ARGUMENT FROM SOCIALITY

I have argued in the previous chapter that the participation of the material world in the process of knowledge production does not suffice to justify experimenters' confidence in experimental results. The participation of the material world in knowledge production may be seriously limited by the constraining effect of scientific culture and by scientists' actions during the process of knowledge production. In this chapter I argue that the justification of the epistemic value of experimental results requires greater autonomy from both, scientific culture and scientists' predispositions. I then show that this autonomy is acquired via the exercise of *collective criticism*, that is, the critical examination of the results of practice by those who did not take part in the production process.

The scrutiny of scientific results from the viewpoint of other scientific cultures and beliefs promotes the autonomy of the process of knowledge production from them both. The confrontation of a given problem-solution, which was generated within the framework of a particular set of items of culture and beliefs, with other cultures and beliefs gives rise to new problem-situations that extend the knowledge production process. Previously neglected courses of action are then identified and explored. In the process, the constraints established beliefs create are gradually removed. *The epistemic import of the social dimension of knowledge production is that the results that eventually come to be accepted by the wider community of scientists are not only justified but they have also a high potential for generating new knowledge.* The results that win collective assent by the relevant community of scientists are supported by robust three-way coherences obtained with the articulation of a wider set of items of culture and beliefs. These robust three-way coherences have a high potential to generate new knowledge that is generated by the resolution of new problems the collective work brought about. Of course, as will become apparent below, the social dimension of knowledge may also create obstacles to the growth of knowledge. In particular, it may lead to conformity, opportunism and dogmatism. These obstacles may then be counteracted by a social organisation of science that promotes *effective collective criticism*.

THE NEGLECT OF THE SOCIAL DIMENSION

Even though the social dimension of knowledge production is pervasive in the studies of experiment, its epistemic value was not explicitly noticed. This is not surprising given that these were intended as responses to the radical reading of the sociological accounts of experiment. These studies have instead emphasised the role of the material world in knowledge production. The collective nature of knowledge production appeared only implicitly in the form of the socially ratified items of scientific culture, the 'resources of practice', the 'epistemological strategies' or the 'constraints' of experiment (notwithstanding the different views regarding the actual role these items play). I will now look in more detail how Pickering, Hacking, Franklin and Galison have dealt with the social dimension of scientific experimentation. The account of David Gooding is then introduced to better illustrate the epistemic value of the collective aspect of the experimental process of knowledge production.

Pickering's account of experiment focuses on the ways by which experimenters engage the material world in knowledge production. The collective aspect of knowledge production appears only as a means to solve the incommensurability of the experimental systems. On his view, experimental systems realise particular engagements with the material world. Each of which depends on the particular actions and reactions of the material world. As a result, they realise "specific captures and framings" of material agency. In his own words:

The representational chains of science terminate not in "the world itself" but in specific captures and framings of material agency. And there seems to be no reason to suppose there exists any particular especially privileged – "true" – machinic mapping of the world. To the contrary, one can easily imagine that an indefinite variety of machinic grips on the world are attainable, attachable in all sorts of ways to all sorts of representational chains. Here, therefore, we glimpse the possibility of a form of incommensurability that cannot be articulated within the representational idiom, an incommensurability in captures and framings of material agency and in the representational chains that terminate in them. This is the new way in which the mangle makes incommensurability thinkable. (Pickering 1995b, p. 188)

In the presence of conflicting results, no objective resolution is forthcoming (see section 2.9). There will be no shared standards for the appraisal of experimental results. If at all, a solution is driven by factors external to the alternatives in dispute. It will most likely, on Pickering's view, be driven by conformity to dominant and established scientific culture. Thus, scientific communities while affecting the knowledge production process do so in a non-rational way.

As we have seen in the previous chapter, Hacking highlights the role of both the established items of scientific culture and the agency of the material world. And similarly to Pickering, he stresses the incommensurability of experimental systems. However, according to Hacking the joint effect of scientific culture and material agency introduces deterministic component to the results of experiment that undercuts the social dimension of knowledge production. Knowledge is whatever results from the current items of scientific culture. This point is most explicit in his *The Social Construction of What?* where Hacking argues that while science as a process is undoubtedly a social activity, it is trivially so. Science as a product, that is, as a set of accumulated knowledge, does not need to invoke its context of production (1999, p. 67). In short, Hacking dismisses any role for the collective aspect of knowledge production, positive or negative.

Franklin, too, rejects any social basis for the validation of beliefs. In fact, the 'epistemological strategies' of experiment (cf. section 2.3) are taken as socially neutral arguments that provide reason for rational belief because they can be independently justified, i.e. they are autonomous from any social consideration. On his view, insofar as 'independent' and 'reasonable' arguments can be provided that justify the 'rationality' of experimental practice, there is no room for considering the epistemic value of the social dimension of knowledge production (1989, p. 459). Of course, Franklin's presupposition is that all relevant strategies are applied and stretched to their limits. But that scientists do not fully explore their problem-situations is

precisely the kernel of the epistemic difficulties informing the experimental process of knowledge production this thesis addresses. From the accounts reviewed in Chapter 2, Galison's is the most socially loaded. For this reason Galison's views are discussed in more detail below.

THE HEURISTIC ROLE OF SCIENTIFIC CULTURE

Galison highlights the heuristic role of scientific culture in helping scientists to bring an end to the experimental process of knowledge production (cf. section 2.7). It does that by guiding and reducing the courses of action available to them. Even though he does not explain how the items of scientific culture come into being and how they are reinforced in practice, it is assumed that scientists accept them as a secure basis for knowledge production. This view, however, is suitable to a Kuhnian reading. Scientists are educated within a disciplinary culture that provides them with a set of tools for their practice. This training contributes to the relative stability of disciplinary fields insofar as scientists socialised in a particular field do not easily transfer their explicit and tacit knowledge to other fields. Scientific training therefore imposes on scientists a cognitive limitation preventing the development of work outside their fields. Scientific institutions, in turn, enforce established standards to ensure that good work is done in their specific fields. By controlling scientists' practices, the institutions of science warrant the objectivity of the scientific enterprise. The detrimental effect of this control is not identified. Scientific culture constrains but this does not seem to pose any significant problem to scientific discovery.

Galison's case studies are very illuminating in depicting how the constraining role of scientific culture may be actually alleviated by the collective process of knowledge production. The conceptualisation of experimentation as an error-elimination endeavour that lacks a logical termination point presupposes a severe social process of evaluation. The absence of a universal criterion that indicates the end of experiments renders experimental results very vulnerable to public scrutiny. As a result, scientists are compelled to generate and accumulate a significant amount of persuasive arguments that stand up to critical challenges. Even though this line of reasoning is not explicit, it does come out very clearly in Galison's reconstruction of experimental physics of late XIX century and XX century. In fact, Galison's *How Experiments End* strategically focuses on disputes between rival research groups to highlight "the characteristic ways each constructed a persuasive demonstration" (1987, p. 14). He presents varied episodes that illustrate how scientists develop competing solutions to a problem, and how they revise and correct these solutions as a result of their confrontation.

Contra Pickering, Galison provides examples where dialogue was possible among groups of different experimental and theoretical traditions and consensus achieved. But he notes that the degree of persuasiveness of the arguments varied between the groups. This is how Galison synthesises the dispute concerning the existence of neutral currents between the group of scientists working at CERN, Geneva, and the group located at FNAL, Illinois:

One group arrived at a temporary conclusion in sharp disagreement with later accepted results. For several weeks some of the researchers were persuaded that their experiment provided evidence that neutral currents did not exist at the level predicted by the Glashow-Weinberg-Salam theory. The discrepancy with theory and with their competitors forced the team to go back over their procedures, producing an almost day-to-day record of the struggle to sort signal from background... each working subgroup assembles its own style of argument... Some of the experienced bubble-chamber experimentalists, for example, came to trust "golden events" more than statistical demonstrations using computer simulations. Those experimentalists more habituated to electronic detectors felt otherwise. (op. cit., p. 18-9)

Thus, consensus does not require that the members involved in a dispute agree on all the arguments put forward. It requires instead the accumulation of enough arguments so that particular scientists find a subset of these sufficient to argue for the validity of the result. Insofar as all objections are met, the community as a whole can consent on the final result. On this view, then, consensus is based on a composite argument that meets all charges levelled against the experimental result. In more abstract terms:

An argument develops within the community as a whole that is of this form: We think there is an object a, because in our historical period there are only a finite set of imitative effects that are plausible: b, c, d, e. Furthermore, we have shown that our phenomenon is not one of these. Importantly, *which alternatives are taken seriously at a given period depends on the prior instrumental and theoretical commitments of the experimenters.* (op. cit., p. 132 emphasis added)

Galison's historical account of experimental physics clearly shows how the engagement of different communities of researchers in the appraisal of conflicting solutions promotes the exploration of unforeseen possibilities. The exploration of these possibilities eventually resolves the ongoing conflict. When arrived at, the consensus on one of the competing solutions is not only more reliable, but it also carries new knowledge. Thus, collectively accepted problem-solutions are epistemically stronger because they are produced by a process that raises relevant questions for solving the problem at hand. In addition, these problem-solutions are supported by robust three-way coherences established with a wider set of cultures and beliefs.

INTERACTING WITH THE SOCIAL WORLD

David Gooding's study of experiment is the most explicit in recognising the epistemic role of the social dimension of knowledge production. Gooding asserts that not only does knowledge production require human intervention in the natural world, but that it also requires social interaction with other scientists. The motto that pervades his account is "observation implicates both the natural as well as the social world". The underlying presupposition is that "the natural world does not constrain any one observer sufficiently, *independently* of any other, to explain observational consensus. Observation implicates other observers as well as the natural world" (1990, p. 76, emphasis in original). The point is that individual causal interaction

with the natural world is not sufficient to warrant self-assurance of what is observed.³⁷ This is particularly the case when observers face novel phenomenon that they try to make sense of “as they go along”. Under these circumstances, different perceptions must be tried against each other until the observers reach agreement about what is being experienced.

Gooding emphasises that consensus is *not* a matter of certifying that each individual has the same experience when he has independent perceptual access to the same object. It arises instead from a social process of consultation based on the contribution each individual brings to the process of knowledge production. The implication of this is that scientists’ interaction with each other carries epistemic weight, as does the interaction with the material world, an interaction that is also part of the construction of the novel claims to knowledge.

The collective dimension of knowledge production already occurs at the private stage of knowledge production, that is, prior to the public report of the experimental results. At this stage, knowledge production is first and foremost a cognitive process whereby scientists jointly make sense of their practices and outcomes. This collective process revolves around the formulation, exchange and revision of tentative interpretations of the novel phenomenal input or the ‘construals’ of phenomena. As Gooding put it, this is a process in which “observers exchange tentative constructs or *construals* of their personal experience”... [and] construe and reconstrue their own experience in the light of what other observers take theirs to be” (op. cit., p. 23, emphasis in original). Thus, already at the private stage, the collective aspect of knowledge production expands the production process because different scientists have different interpretations of their practices and outcomes.

The necessity of submitting the results of science to public scrutiny also affects the private stages of knowledge production. Internal interaction is influenced by the subsequent stage of public scrutiny insofar as scientists know that their results are to be assessed by others. Because claims to knowledge are to be subjected to the critical scrutiny of a wider collective of scientists that extend beyond those directly involved in the production process, knowledge producers already take into account other viewpoints in the generation of their results.³⁸ However, it is in the public stage of knowledge production where problem-solutions are actually put to the test.

The differentiated epistemic value between the private and public stages of the process of collective criticism is expressed by the fact that any claim to knowledge must win the consent of the wider community of scientists. The publicity of a coherent result does not warrant its immediate acceptance. It is only when the wider

community assents to the experimental results that these are acknowledged as scientific results. As Gooding has it:

[T]he epistemological warrant of experiment resides outside any particular experiment or its outcomes... empirical judgements necessarily have a consensual basis. The factual status of the result of an experiment – and therefore its epistemic force – lies in evaluative judgements made by groups of experts about the quality of experiment and its theoretical plausibility (both as to method and to outcomes). On these judgements in turn depend the ‘externality’ of phenomena or data, that is, their agreed status as real, self-evident, possible, or as artifacts... the ‘correct outcome’ is [thereby] socially validated. (1990, p. 211)

It is now clearer how the social dimension of knowledge production enhances the epistemic value of experimental results. The collective acceptance of experimental results furnishes further assurance that these results are the outcome of a reliable process of knowledge production. It ratifies that these results are produced by a properly functioning material apparatus that can yield information about the material world. That is, it assures that the results are produced by the agency of the material world rather than by a faulty material apparatus. The private results then turn into self-evident natural facts detached from the contexts wherein they were generated:

A natural phenomenon is at first identified with its discoverer and a few initiates and would not exist without them. To become accepted as part of scientists’ experience that phenomenon must be transferred from the personal realm to the public domain, where it can be witnessed by all. To be seen as a natural fact it must be seen to exist independently of any particular person, laboratory or experimental technique. (Gooding 1985, p. 105)

According to Gooding the public stage of knowledge production is highly dependent on the experimenters’ ability in communicating their findings to lay observers.³⁹ Successful communication, in turn, depends on experimenters’ representational skills that allow other observers to witness phenomena without having to “tread every turn of the path travelled to produce them”. Successful representations enable lay observers to ‘witness’ the phenomena and thereby persuade that “they could in practice reproduce the same process and would get the same correspondence between concepts and precepts” (1990, p. 167). This is, in fact, what allows the public evaluation of both the processes and products of experiment.

These representations are simplified and idealised versions of practice and phenomena. They erase as far as possible the participation of the scientists in the process of knowledge production, detaching the experimental phenomena from the actions and the procedures necessary to generate them. The experimental process is then reconstructed as following a single and linear sequence of operations in which every step leads directly to the single and unique outcome conveyed by the output

³⁷ Gooding speaks of ‘observation’ and not of experiment to stress the fact that even the exercise of visual perception requires agency in the material world – ‘it is action that enables seeing’ – and in the social world to gain intelligibility of what is seen.

³⁸ Mäki introduces the notion of ‘gate-keeping conditions and procedures’ which are “tacitly or explicitly accepted by a disciplinary profession that delimit the protagonists and the contents and styles of claims and arguments to those that are recognized by the profession as worthy of further consideration” (1993, p. 96). Besides the conformity to established culture, in order to be recognised as competent, Mäki suggests, scientists exercise a sort of ‘private rhetoric’ directed to an ‘internal audience’. This rhetoric is an internal kind testing, performed in the mind of scientists, before ideas and arguments are presented to an external audience (ibid, p. 97).

³⁹ A lay observer is “anyone not yet familiar with the observational practices generally necessary to produce and see a phenomenon. Most scientists are lay observers with respect to new phenomena” (1985, p.133, n. 8).

phenomenal model. Only the steps required to produce the phenomenal output are presented. Dead-ends, draw-backs, failed trials and errors are all discarded from public reports. Gooding notes that this detachment has a significant contribution to the 'natural' status of the phenomena depicted in the representations.⁴⁰

[I]n the retrospective accounts we find a more direct interplay between changes in apparatus or models, empirical outcomes, and propositions explaining the purpose of such changes and the significance of their results. The interplay is more dialectical... in that there is a more even matching of outcomes and responses... it gives a more structured appearance to the course of... research. (Gooding 1990, p. 169)

According to Gooding the public reports render experiments 'transparent' in the sense that "the apparatus and procedures appear to contribute nothing to what the experiment shows" (1985, p. 107). This means that the published reports of experiments constitute an important post-experimental part of practice intended to persuade the community of scientists that the experimental results are not artefacts of the experimental procedure. And, in particular, they aim at persuading that the phenomena that were made accessible only through scientists' agency are accepted as independent of that activity. As Gooding put it, in the public accounts scientists' agency becomes invisible and the phenomena appear as a 'residue' of experimentation independent of human intervention (1989, p. 217).⁴¹

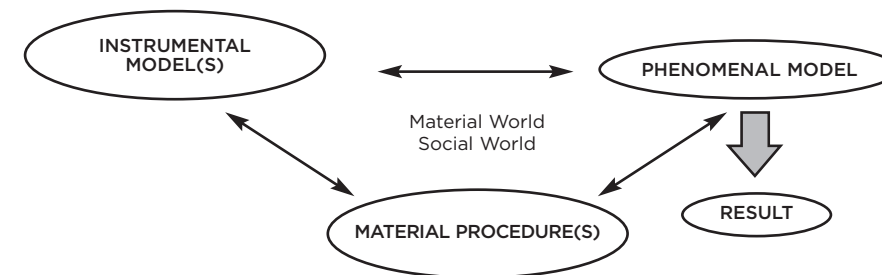
To conclude, Gooding's point is that experimentation is a collective enterprise and that the collective nature of this enterprise has epistemic value. Scientists undertake joint endeavours and submit their outcomes to the critical scrutiny of the wider community. The epistemic role of the social dimension of knowledge production depends on how far the collective process of criticism is carried out in actual practice. This is the point I now turn to.

THE ARGUMENT FROM SOCIALITY

The *argument from sociality* highlights the contribution of the social dimension of knowledge production to the generation of epistemically sound novel claims to knowledge. The fact that knowledge production is the joint endeavour of a collective of scientists, who may be educated in various cultures and may be prone to different commitments, renders the results of science less vulnerable to the constraining effects of particular items of scientific culture and prior beliefs. To put it in Pickering's terms, interaction with other scientists causes 'resistances' to human intentions that must be 'accommodated' if scientists want the larger scientific community to accept their results as products of science. These are 'social' resistances that stem from the 'agency' of subsets

of scientific culture and beliefs and that give rise to alternative interpretations of the instrumental and the phenomenal models and thereby extend scientists' intervention in the material world. In the process, the range of admissible interpretations narrows down and consensus is reached. When consensus is arrived at, the community of scientists assents on the interpretation of practice as conveyed by the consensual versions of the instrumental and the phenomenal models. On this view, then, consensus is a social outcome that obtains after confronting different sets of items of culture and beliefs that enhances the autonomy of the results of science from any particular subset of them.⁴² In figure 3.1, the large-sized shape of the material procedures and the instrumental models and the small-sized shapes of the phenomenal model and experimental result capture this intuition.

FIG. 3.1 The Social World and the three-way experimental coherence



The dependence of knowledge production on extant culture and upheld beliefs are never fully eradicated. Nor should they be so. As explained in the previous chapter, in science there is no such thing as starting from scratch. The point is that the social dimension of knowledge production can overcome, to some extent, such dependence. This role depends, however, on the composition of the community of scientists and on scientists' predisposition to put established beliefs to the test. The epistemic value of the social dimension of knowledge production, or the extension and the scope of the knowledge production process, hinges on the level of heterogeneity of the cultures and of the beliefs of the scientists participating in the production process. The more heterogeneous the communities, the more fruitful problem-situations are likely to be identified and relevant courses of action explored. To put it bluntly, *whereas the epistemic role of the materiality of experiment may be limited by the plasticity of the experimental systems, the epistemic role of the sociality of*

⁴⁰ To a certain extent this explains why experimentation has been considered a non-problematic activity. Experimental narratives have been erroneously taken as descriptions of experimental practice rather than reconstructions.

⁴¹ This point had already been remarked by Latour and Woolgar (1986[1979]) as a necessary stage of the process of transforming experimental results into 'scientific facts'.

⁴² In experimental economics, Timo Tammi has explored the social dimension of knowledge production along similar lines. In (1999a) he shows that the cooperative work of theoretical and experimental economists can produce consensual results notwithstanding divergences about the significance and role of the experimental method of economics. In (1999b) Tammi shows that the larger collective of experimenters do try to work out incompatible results and thereby improve coherence among extant products of science. (These topics will be explored at length in the second part of this thesis)

experiment may be limited by the homogeneity of the scientific community. This is not to say that any scientific culture and set of beliefs can and should be brought into the process of knowledge production. The analysis of the sociality of the process of knowledge production must take into account the relevant kind and the adequate level of heterogeneity for the problem-situation at hand. This topic is, however, postponed until section 3.6. Before that, the specific contours of the epistemic value of the social dimension of knowledge production need further scrutiny.

To recapitulate, the *argument from coherence* asserts that the three-way coherence achieved among the three components of the experimental system provides scientists with reasons for believing that they have made sense of their practice and resulting outcomes. The *argument from materiality* reinforces the epistemic value of achieved coherences by noting the role of the material world in promoting the revision of scientific culture and that of scientists' prior beliefs. The *argument from sociality* bestows added epistemic value to the three-way coherence by highlighting the role of the collective aspect of knowledge production in bringing a larger number of items of scientific culture and beliefs to the production process. In sum:

The social dimension of knowledge production promotes the participation of a vast number of varied items of scientific culture and beliefs in the production process. These bring about alternative interpretations of experimental practice and the results of practice which improves the potential for the generation of epistemically sound novel claims to knowledge.

The three arguments in conjunction support the main result of SEE:

The greater the number and the wider the heterogeneity of the resources (material, conceptual and social) involved in knowledge production, the higher the epistemic status of the three-way coherence and the experimental result it supports. The solution arrived at is epistemically sound because it is the result of practices that have explored relevant courses of action to finding a solution to the problem at hand.

It should be noted that heterogeneity is neither a sufficient nor is it a necessary condition for the production of epistemically sound results. The point is simply that having recourse to heterogeneous resources is epistemically significant when experimenting with plastic experimental systems to produce novel phenomena.

The social validation of scientific knowledge is not an original idea. It has been the focus of the sociology of science and the sociology of scientific knowledge (notwithstanding the philosophical implications derived) and it has recently attracted the interest of the philosophers of science (as will become apparent below). SEE highlights how the collective nature of knowledge production contributes both to the discovery and the validation of experimental results. It takes as a non-negligible part of experimental practice the exchange and discussion of scientific results. This is what scientists do when they engage in internal discussions with their peers, when they present their results in conferences, when they write papers and submit them to publication in credible journals, and when they engage in critical debates by challenging others' results and responding to charges targeted at them.

The prominence and relevance of scientific debate have inspired a recent interest in the study of the rhetoric of science.⁴³ This exercise has been depicted, not surprisingly, as a matter of persuading one's audience where persuasion is measured by the success in transforming the audience's beliefs. From what has been argued thus far neither should it be surprising that the effectiveness of the persuasive arguments is taken to be a function of the relations of coherence these arguments establish with upheld beliefs (cf. Mäki 1995, p.1304-5). The next section illustrates this with another episode of the hunting of the quark.

3.3 THE HUNTING OF QUARKS 3: THE SOCIAL RESOLUTION

In chapter 2 two sets of experiments were presented that yielded conflicting results. This was not presented as a demonstration of misconduct on the part of the two teams of researchers, or any single one of them for that matter. Both groups of researchers attempted to generate credible results that could win collective assent. But because the Stanford experiments conflicted with established knowledge, they were subjected to more intense scrutiny than the others. In fact, at the time the Stanford team announced their results, Pickering says, "it was clear that any claim to have found experimental evidence for isolated quarks would be viewed against a hostile background of previous unsuccessful experimental searches, compounded by theoretical accommodation to this lack of success" (1981, p. 217). This hostility provides evidence for the assertion that scientific communities do resist results that conflict with their culture. However, like material resistances, social resistances are not insurmountable. They can be accommodated in subsequent practice as scientists satisfactorily respond to critical challenges. In the process, more items of culture and beliefs are brought to knowledge production which promote the generation of sound novel claims to knowledge. This was indeed the case with the quark experiments.

The 'social resistance' encountered led the Stanford team to improve their experimental equipment and argumentation in support of the phenomenal model that accounted for the existence of quarks (cf. section 2.9). But the group did more than that. It put forward the 'tungsten hypothesis' which organised the conflicting results. The tungsten hypothesis aimed at responding to a particular critique. In one of the experiments one ball measured in one run a charge of $-1/3e$, and in a subsequent run it carried a charge of 0. To the Fairbank team it suggested that the ball in the first run had a quark that was lost between runs. But the critics interpreted the result of the first run as being caused by some spurious effect. Having observed that the balls that carried fractional charges were those that had been supported on a tungsten dish in the heat-treatment procedure, the Stanford group conjectured that quarks could be associated with the tungsten of the substrate.

⁴³ See (Mäki, 1995) for references in various disciplines. In economics, the seminal contribution is McCloskey's (1985) *The Rhetoric of Economics*.

Under this reading, the quarks would be transferred to the niobium balls during heat treatment, and then lost between runs. The presence of quarks in the first run and their absence in subsequent runs thus became plausible by suggesting that quarks were not uniformly distributed within the ball and that they could be dislodged from the surface of the ball between runs. This interpretation rendered the loss of charge intelligible and thereby supported the hypothesis that quarks exist. But more crucially, it also solved the conflict with other experiments. Tungsten had only been used in the Stanford experiments.

The response of the Stanford team provides further evidence for the ease with which highly plastic systems re-establish destabilised relations of coherence. But when taking into account the social dimension of knowledge production, plastic experimental systems are not as problematic. The articulation of conflicting results captured the attention and interest of the parties in the dispute and, as a result, guaranteed the continuation of the production process. By organising the conflicting results of the research program in a unifying account, the tungsten hypothesis attracted the interest of high-energy theoretical physicists who became receptive to the hypothesis of isolated quarks. New phenomenal models were developed which further elaborated the tungsten hypothesis and improved the articulation of the Stanford results with those of the quark community.

The theoretical developments suggested experimental work that could be easily performed with the conventional Millikan apparatus. The tungsten hypothesis implied a greater abundance of quarks in tungsten which should be detectable by the Millikan apparatus. However, no evidence was found for fractional charges.⁴⁴ As a result, the tungsten hypothesis was abandoned and the loss of charge from the niobium balls became again under suspicion. The results obtained were deemed conclusive because the tungsten experiments could have recourse to the standard Millikan apparatus. The use of a well-established and understood apparatus gave no reasons for doubting the results. Had the results been positive, the quark program would probably have continued.

This episode of the quark program illustrates how the practice of the Stanford team was affected by the wider social context and how the community reacted to the team's responses. Fairbank and his team wanted their results to be accepted and wanted to be recognised as competent experimenters. To that end they made an extra effort in articulating their controversial results with those of others. The strategy turned out to be successful. By providing a coherent interpretation of the conflicting results, the Fairbank team managed to attract the community's interest. New problem-situations emerged and new experiments were carried out. But it is important to note that these problem-situations arose because the community was composed of scientists trained in different cultures. As a result, they could conceive new ways of extending the production process. As more scientists engaged in the process, more disparate resources of practice and commitments were brought to the attainment of a common goal – the search of quarks. As the research evolved and involved a wider number of scientists, the

results of practice became more conclusive. The accumulation of negative results for an already weak claim determined the conclusion of the research program.

At this juncture, it is interesting to note that in cases where the debate has already implicated a wide community of researchers, the generation of supporting evidence to a controversial claim may prolong the research program whereas the generation of one additional negative result may have the epistemic force of an *experimentum crucis* and dictate the conclusion of a research program. More importantly, however, the analysis of the social process of production of experimental results shows that allegedly incommensurable experiments can in fact be confronted with each other and thereby the subcultures and commitments of the researchers involved. *The Hunting of the Quark* also suggests that the cooperative effort of the scientists hinges upon three conditions:

1. *The producers of controversial results must bear the burden of articulating the results obtained with established knowledge so as to render them significant to other scientists.*
2. *The scientific community must be receptive to novel and unexpected results and open to engage in the solving problems prompted by the research of other scientists.*
3. *The realisation of these two conditions, in turn, requires an adequate organisation of scientific production that encourages and facilitates both kinds of behaviour.*

3.4 THE SOCIAL DIMENSION IN THE PHILOSOPHY OF SCIENCE

The philosophers of experiment have not yet addressed the social dimension of the experimental process of knowledge production. Symptomatic of this state of affairs is the recently edited book by Hans Radder (2003) devoted to “The Philosophy of Scientific Experimentation” in which none of the thirteen contributions took up the topic. This can be easily ascertained by the following synthesis of one reviewer:

Each essay examines one of or more of six interconnected themes that run through the book: (1) the philosophical implications of actively and intentionally interfering with the material world while conducting experiments; (2) issues of causality in experiments; (3) the link between science and technology; (4) the role of theory in experimentation; (5) the impact of modelling and computer simulation on experimentation; (6) the philosophical implications of the design, operation and use of scientific instruments. (Boon 2004, p. 221)

The ‘philosophers of theory’, in contrast, have recently recognised the epistemic role of the social aspect of knowledge production. Indeed, this recognition gave rise to the new branch of philosophy of science, social epistemology. Even though these contributions do not specifically address experimentation, they do offer important insights to the social epistemology of experiment under construction here. Some earlier inputs are first looked at, however.

⁴⁴ Pickering quotes four differently co-authored articles (1981, p. 232, n. 32).

EARLY CONTRIBUTIONS TO THE SOCIAL EPISTEMOLOGY

Philosophers have already dealt with the social aspect of epistemic justification. The problem of ‘testimony’, i.e. the problem of justifying acceptance of assertions made by others, has concerned philosophers such as Plato, John Locke and David Hume. However, the systematic study of the social dimension of the processes of knowledge production in science and its role in epistemic justification is a recent endeavour. Until then epistemology has been ‘individualistic’ in that the epistemic justification have focused on the cognitive processes of individuals in the acquisition of justified, rational or true beliefs.

The ‘social’ field of epistemology was first demarcated by Alvin Goldman (1978, 1986, 1987) with his subdivision of epistemology into two branches: individual epistemology and social epistemology. On his account, individual epistemology concerns the psychological processes of the epistemic subject, while social epistemology concerns the social processes whereby the epistemic subjects interact with others who influence their beliefs.

Karl Popper can be considered the modern precursor of social epistemology given the role criticism plays in his account of science. Popper explicitly acknowledges that the collective dimension of knowledge production improves the ‘rational and critical attitude’ of the individual scientists. Popper argues that a Crusonian scientist could never produce scientific results, regardless of his critical effort. The process of knowledge production would always fall short of the collective stage of evaluation. He says:

For there is nobody but himself to check his results; nobody to correct those *prejudices* which are an unavoidable consequence of his peculiar mental history; nobody to help him to get rid of that *strange blindness concerning the inherent possibilities of our own results* which is a consequence of the fact that most of them are reached through comparatively irrelevant approaches. And concerning his scientific papers, it is only in attempts to explain his work to somebody who has not done it that he can acquire the discipline of clear and reasoned communication which too is part of scientific method (1979[1972], p. 219, original emphasis omitted and new emphasis added).

According to Popper the collective process of evaluation enhances the critical ability of the individual scientists in problem-solving insofar as the problem-solutions ought to be communicated to and appraised by others. In addition, it allows the identification and correction of error unperceived by the scientists directly engaged in the resolution of a specific problem-situation. In this regard the social process of evaluation amounts to a “friendly-hostile co-operation of many scientists” (op. cit, p. 217) whereby scientists identify and eliminate each other’s errors. Popper notes that this friendly-hostile co-operation requires adequate social institutions that foster critical discussion. And that it is upon the functioning of these institutions where “all progress, scientific, technological, and political, ultimately depends” (op. cit, p. 218). However, Popper does not say how these social institutions are to contribute to the cooperative interaction among scientists and thereby to the critical attitude of individual scientists.⁴⁵

Popper notes that the social process of evaluation is the more effective the wider the disparity of the views in confrontation. This is the case because the higher will be the potential for learning. More difficult and unforeseen questions can be raised and given due attention. And the effectiveness of critical evaluation is assessed by the extent to which prior beliefs were transformed during that process. In his words:

I think that we may say of a discussion that it was the more fruitful the more the participants were able to learn from it. And this means: the more interesting questions and difficult questions they were asked, the more new answers they were induced to think of, the more they were shaken in their opinions, and the more they could see things differently after the discussion – in short, the more their intellectual horizons were extended. Fruitfulness in this sense will almost always depend on the original gap between the opinions of the participants in the discussion. The greater the gap, the more fruitful the discussion can be. (Popper 1994, p. 35-6)

But Popper’s view of science is above all individualistic. The most significant problems in knowledge production are of a psychological nature. They result from scientists’ predisposition to uphold ingrained beliefs. Consequently, the solutions are also individualistic. They appeal to good conduct and the imposition of restrictions on the available actions of scientists. Indeed, Popper puts great emphasis on how scientists ought to be good scientists. In general, scientists ought to be self-critical in what they do. And in particular, they should observe three norms of conduct. First, the individual scientists should be committed to “do good work in their own particular fields”. Second, they should have a broad range of interests so as to “shun the danger of narrow specialization”. And third, they should divulge their work to non-experts (1994, p. 109-10). And in regard to their field of specialisation, they should follow the critical method – falsificationism – the method of submitting bold theories to severing testing (1959[1934], 1965[1963]). The social dimension of scientific practice is a means to the same effect: to improve the critical attitude of scientists and correct it when it fails.

Paul Feyerabend (1993[1975], 1978) is instead concerned with the detrimental effect of the social conditioning of science. In particular, he is concerned with the inhibiting effect of dominant cultures on the development of alternative viewpoints.⁴⁶ Following John Stuart Mill (1859), Feyerabend advocates the necessity of an unimpeded opportunity for the exercise of the critical discussion of ideas. This is the only way, on his view, by which partiality of belief or opinion can be avoided. The worry is the dogmatic acceptance of beliefs and their crystallisation as ‘myths’. Thus, unlike Kuhn, Feyerabend thinks that different theories, worldviews, cultures, can be confronted with each other. Not only can they be confronted, they should be so

⁴⁵ The organisation of scientific institutions is, to some Popperians, the central message of Popper’s work, and recently it has received the recent attention of Agassi, Bartley, and Radnitzky (see Hands, 2001). As will become apparent below, this topic is also central to social epistemology.

⁴⁶ Mäki distinguishes three conceptions of social conditioning of science (which are not meant to constitute an exhaustive classification): 1) the influence exercised over the goals of scientists’ actions; 2) the justification of scientific results; and 3) the contents of science (1993, p. 82). According to Feyerabend dominant cultures have a negative effect on each one of these dimensions by restricting scientific practice to a narrow domain and the content of science to conformity with established beliefs.

confronted. The opposition of incompatible belief systems is the only way whereby prejudices can be identified.

A scientist who wishes to maximize the empirical content of the views he holds and who wants to understand them as clearly as he possibly can must therefore introduce other views; that is, he must adopt a *pluralistic methodology*. He must compare ideas with other ideas rather than with ‘experience’ and he must try to improve rather than discard the views that have failed in the competition. (1993[1975], p. 21 emphasis in original)⁴⁷

Rather than working within an existing belief system until it gets into difficulties, scientists should propose, develop and tenaciously defend alternative systems to test the dominant system. The confrontation of a given theory with alternative theories forces “the others into greater articulation and all of them contributing, via this process of competition, to the development of our consciousness” (ibid).

On this view, only the confrontation of different cultures can attenuate their constraining effect. This demands from the individual scientists willingness to confront their results with other viewpoints and from the science institutions it requires receptiveness to all forms of culture. Thus, Feyerabend not only acknowledges the negative effect of dominant scientific culture in the production of scientific knowledge in addition to the cognitive limitations of the scientists, but he also believes that the confrontation of different cultures is possible and necessary. To that end, scientists should be allowed to develop alternative viewpoints.

THE SOCIAL EPISTEMOLOGY OF THEORY

The social branch of epistemology attempts to reconcile the sociological explanation of scientists’ beliefs, based on social factors such as professional interests and political ideologies, with the idea that science is a rational endeavour. This is explicitly expressed by Goldman in his reply to Larry Laudan (1977) and William Newton-Smith’s (1981), who took the field of sociology of science as the field of research that concern the study of the beliefs that cannot be given a rational justification. Philosophy, in contrast, is the field that studies scientific rationality.

There is no reason why one and the same belief, scientific or otherwise, cannot be explained both in rational and in social terms. This leaves very wide scope, in principle, for social explanations of belief. I agree with Laudan that it is very unlikely that there are laws governing belief that feature *exclusively* sociological variables in their antecedents. Social or sociological factors could not provide a ‘complete’ explanation of belief. But there is no reason why social factors cannot influence belief-formation in conjunction with, indisputably rational, factors... I would go further, however, and point out that rationality might partly consist in certain forms of social interchange. Some styles of debate and mutual criticism, some modes of doxastic response to the

arguments of others, may be partly constitutive of rationality. So there is no tension between social and rational belief causation. (Goldman, 1987, p. 110-1 emphasis in original)

Subsequent accounts took a step forward. Not only do they argue that the social character of science is compatible with rationality, but also that the social organisation of science may promote compatibility between scientists’ idiosyncrasies and the epistemic goals of science. The position of Kitcher (1990, 1993, 1994) and of Goldman and Shaked (1991) are representative of that. The main thrust of their arguments is that individual attributes (such as objectivity, impartiality, or rationality) are not necessary for the attainment of the goals of science. Nor is strict conformity to the rules of scientific method(s) required to that end. Scientific rationality can be promoted by the right institutional arrangements that make the best epistemic use of the cognitive and psychological traits of individual scientists to further the progress of science (that is, the generation of justified or true beliefs). The central issue to social epistemology is then the study of the social conditions that take advantage of individual attributes, including the different “dispositions to acquire true beliefs or to be moved by reliable belief-forming processes” which may be “surely traceable to the distortions introduced by social biases or personal prejudices” (Kitcher 1994, p. 124). On this account, *diversity* and *competition* are the main criteria around which science should be organised. Scientific communities should encourage the division of cognitive labour among scientists, rather than the pursuit of what may seem the most promising strategy. Division of cognitive labour, so the argument goes, favours the production of theories with different epistemic value and thereby raises the chances of generating justified (true) beliefs. Competition, in turn, ensures that consensus is reached on the best theory, that is, the most warranted theory.

To Kitcher consensus is easily achieved. The competition between rival theories ensures that the proposal that gathers best arguments is the one that receives collective consent. Kitcher demonstrates this by means of an economic model (1993, Ch. 8). He shows that even in a community of self-interested reward-maximising scientists (rather than truth-seeking scientists) occurs a distribution of effort among diverse courses of action. The competition for the monetary reward in such community then forces scientists to improve their standings in response to critical challenges until they accept the best theory. When extending the results of the model to science, Kitcher concedes that there might be non-epistemic goals in operation, but he claims that “the argument emerging when the debate crystallizes should be the *principal* source of power – in other words, that in a competition between the social factors and arguments leading in a contrary direction, the acquisition of power should be more affected by the arguments. (Social factors may retard a decision but not reverse it.)” (1993, p. 201-2, emphasis in original).

Kitcher argues that under the right social arrangements the detrimental effect of scientists’ motivations can be eliminated or put to science’s service or, in his words, make “good epistemic use of the grubbiest motives” (1993, p. 305). But Kitcher’s solution is highly dependent on two basic assumptions of the game theoretical model – the assumption that self-interested and reward-maximising scientists will distribute their cognitive effort evenly, and the assumption that science can be organised as a

⁴⁷ This view was first elaborated by Feyerabend as part of the methodological principle of “theoretical pluralism” (1965), but it later (1975) became part of his anti-methodological stance to encompass distinct cultures, scientific and non-scientific alike.

competitive market that rewards the epistemically efficient scientists. As argued before, scientific communities and scientists tend to concentrate their cognitive effort on problem-situations that can be solved and appraised with the available and the socially established items of scientific culture.⁴⁸ For this reason, the competitive market can hardly be accepted as a representative institution of science.⁴⁹

Miriam Solomon (1992, 1994a, 1994b, 2001) offers a more pessimistic view of science. She acknowledges that “what matters, normatively speaking, is the distribution of empirical and non-empirical decision vectors across a community of investigators” instead of “the thoughts and decisions of individual scientists”. In this account Solomon does not need to assume that scientists are epistemically rational. Scientists may never agree on the same theory regardless of the available arguments in its favour. Nor do diversity and competition guarantee the best epistemic outcome. Epistemic states depend, instead, on the social, cognitive and motivational factors in presence, or on ‘decision vectors’, as she put it. In other words, Solomon introduces a gap between individual and collective rationality.

Decision vectors comprise any factor that may affect the scientific decision-making process and thereby the direction of scientific practice (hence the geometrical terminology). Two grand types of vectors are distinguished according to their epistemic role. ‘Empirical decision vectors’ have all the epistemic value and “are causes of preference for theories with empirical success”. These include increased empirical content and theoretical consistency with data. ‘Non-empirical decision vectors’ are non-epistemic. They “are other reasons or causes for choice”, which may comprise ‘motivational’ vectors such as pride and conservativeness; ‘social’ vectors such as preference for an ideology; and ‘theoretical’ vectors such as preference for simplicity and elegance (2001, p. 53-8). Even though only the empirical vectors possess epistemic value, any of these factors have an equal chance of influencing the decision-making process in science.

Solomon argues that the non-empirical or the non-epistemic factors are responsible for scientists’ overall resistance to abandoning favoured instruments, methods, standards, and theories and that this resistance is responsible for undesirable epistemic states at the system level. There is therefore no guarantee that the community’s epistemic state is the best given the available knowledge at hand. And conversely, she remarks that adequate epistemic states may be obtained irrespective of scientists’ motivations and cognitive dispositions. For instance, dissent may be justified whenever individual scientists distribute their cognitive effort so as to examine all worthwhile

theories; and consensus may obtain on the best theory even though the underlying decisions are not based on the critical scrutiny of the merits and demerits of the available proposals. In fact, Solomon argues, consensus is reached when the consensual theory performs well on the basis of locally situated standards and interests.

On Solomon’s view, science should be organised so as to implement a distribution of research effort, defined in terms of decision vectors, that contributes most to the ‘maximisation of empirical success’ (op. cit., p. 117). However, both empirical and non-empirical decision vectors are necessary conditions of knowledge production. Whereas the non-empirical decision-vectors are crucial for the development and exploration of competing theories and thus a pre-requisite for the construction of empirically successful theories, the empirical decision-vectors are (or should be) crucial for the selection of the theory that possesses the highest epistemic value. Thus, like Kitcher, Solomon argues that the attainment of the epistemic goals of science require a social organisation of science that promotes diversity so that all worthwhile avenues are explored. However, on Solomon’s view an epistemic state of dissent is normatively appropriate only when the theories on which there is dissent display an *equitable* distribution of empirical decision vectors and an *equal* distribution of non-empirical decision-vectors (op. cit., p. 117-8). That is, the empirical success of the alternatives must be weighed by the magnitude of the empirical accomplishments.⁵⁰ In contrast, non-empirical decision-vectors, which are epistemically empty, can be distributed in equal numbers among the various alternatives. In sum, dissent is an appropriate state whenever the alternative theories in dispute have equivalent empirical successes and an equal share of biasing factors supporting them. Consensus, on the other hand, is normatively appropriate in the special case in which the consensual theory gathers all the empirical decision vectors and all the empirical successes available in its domain of inquiry.

It is not difficult to see that Solomon’s account is too strict. Diversity is adequate only when the alternative theories in dispute have equivalent empirical groundings. Because the development of competing theories is not synchronic, her framework can only be applicable after all proposals have been fully explored and their empirical value assessed. Until then a higher degree of tolerance must be afforded to dissent so that competing theories may gather enough empirical arguments. Conversely, only rarely if at all can consensus be deemed an epistemically valid state if the consensus theory must concentrate all the empirical success.⁵¹

Helen Longino (1990, 1991, 1994, 2002a) also attempts to justify the rationality of science while accommodating the negative effect of the social factors in science. The recognition that neither method nor empirical evidence constitute the sole determinants of knowledge production is Longino’s point of departure. She stresses that background beliefs are also relevant. They determine the extent to which a particular piece of evidence is considered significant to a particular hypothesis. The

⁴⁸ The purpose of this section is to review how the social dimension of knowledge production has been accounted for in the philosophy of science, rather than provide a critical assessment of it. However, it should be clear that this thesis addresses the problems overlooked by Kitcher. It is devoted to the problematic of the lack or potential lack of diversity in science due to both scientists’ cognitive limitations and the obstructing effect of scientific culture. More direct criticisms can, however, be found in Hands (1995), Mirowski (1996) and Solomon (1994b).

⁴⁹ This issue has been recently addressed in the field of the economics of scientific knowledge. In this respect, it is interesting to note that while philosophers of science (as well as other students of science) take the concept of the competitive market as a good model of and for scientific institutions, economists seem to characterise the organisation of the discipline of economics as a monopoly rather than a competitive market and advocate regulatory practices thereof (cf. Mäki 1999). Dasgupta and David (1994) advocate two different modes of organisation of science according to the nature of scientific knowledge: a nonmarket institutional structure for basic science and a competitive market for applied science. For a review of this literature see Hands (2002, ch. 8).

⁵⁰ This magnitude should be appraised in terms of the degree of replicability of the empirical results and the theories’ ability to generate novel unanticipated results (op. cit., p. 28).

⁵¹ As mentioned, the goal here is not to provide a criticism of alternative social epistemologies. However, Solomon’s account raises many questions, such as the election of empirical variables as the sole epistemic factors.

problem is that these background beliefs are particularly permeable to the influence of ideology and social values. In her words:

Because background assumptions can be and most frequently are invisible to the members of the scientific community for which they are background, and because unreflective acceptance of such assumptions can come to define what it is to be a member of such community (thus making criticism impossible), effective criticism of background assumptions requires the presence and expression of alternative points of view. (Longino 1991, p. 670)

Longino is concerned with the appraisal of that part of scientific culture which is unreflectively accepted by the members of a particular scientific community. Objectivity on this view requires the possibility of controlling the influence of subjective preference at the level of these unconsciously held background beliefs. This is achieved by taking into account viewpoints representative of different social perspectives.

Like other social epistemologists, Longino recognises that the social process of evaluation requires special institutions. Ideally, it requires the observation of four criteria. Firstly, “there must be publicly recognized forums for the criticism of evidence, of methods, and of assumptions and reasoning”. This means that the scientific community at large must recognise the role of criticism and be supportive of its practice by creating a public space to that end. But a critical community is not simply one that tolerates dissent. Scientists must actively engage in the practice of criticism, by criticising and responding to counterarguments, and by being open to adjust their views as a result of it. Therefore, the second condition is that “there must be uptake of criticism”. The active engagement of scientists in the critical discussion requires that those who hold the position being criticised perceive the critiques as relevant. And this presupposes the existence of common standards so that the parties involved may agree on the points of convergence and divergence and the means by which the conflict may be resolved. Thus, the third condition is that “there must be publicly recognized standards by reference to which theories, hypotheses, and observational practices are evaluated and by appeal to which criticism is made relevant to the goals of the inquiring community”. Therefore, the confrontation of different scientific cultures requires a common set of standards that must be shared among the parties involved in the critical discussion. If not, the criticisms are not perceived as relevant. Fourthly, there must be “tempered equality” of intellectual authority to guarantee that all relevant perspectives can be represented, though, qualified by the intellectual differences that may exist between them (2002a, p. 129-31).

Even though this is meant as a characterisation of an ‘idealised epistemic community’, i.e. the state scientific communities should aspire to, the four criteria may be taken as general guiding rules to that end. They can be conceived of as norms applying to the social processes of knowledge production to warrant “effective or transformative criticism”, i.e. norms that ensure that the theories and hypotheses accepted by the community do not incorporate the idiosyncratic biases of an individual or group. To that effect, the fourth condition is pivotal. That is, the rationality of science ultimately lies on the promotion of intellectual democracy whereby all participants are given a fair hearing and, in particular, the members of social minorities. Minorities are in an especially advantageous position to identify and criticise the dominant value

system which is unconsciously presupposed by the majority. From this it follows that the possibility of effective criticism requires a social organisation of science that ensures the accessibility of the relevant members to the critical debate, and that takes active steps to ensure that alternative points of view are being developed.

It is now not difficult to see that the possibility of an effective critical scientific community depends on the heterogeneity of the collective of scientists, or the presence of alternative points of view, and on the existence of shared standards for the exercise of criticism. However, the engagement of groups from different communities might make it difficult reaching agreement on the criteria of critical appraisal. As Solomon noted, different communities adhere to different sets of standards which may be perceived as impeding factors to critical discussion, and thereby reasons for dismissing criticisms as relevant. This is even more problematic given that the standards of evaluation are not often neutral to the positions under confrontation. Proponents of a given position tend to favour a sympathetic set of standards whereas the opponents tend to favour one that renders that position more vulnerable to challenge. Failure in reaching agreement on common standards may be disruptive of criticism inasmuch as proponents will not perceive critical charges as relevant and, as a result, will not respond to them. This explains why critical appraisal is easier within subspecialties where a significant part of scientific culture is shared and, consequently, communication is relatively full and professional judgment relatively unanimous. However, the potential for identifying and scrutinising unreflectively held beliefs is narrow.

This points to the importance of having consensual and agreed upon standards for the effective ‘uptake of criticism’. Moreover, as the analysis of the quark research program brought to light, dialogue between different communities requires some work of ‘translation’. That is, the ‘challengers’ of some culture must have the burden of articulating controversial results with established knowledge and eventually suggest interesting and fruitful problem-situations. The ‘defenders’ of extant culture, in turn, must feel compelled to respond to those challenges. These seem to be two additional necessary conditions for the uptake of criticism. The danger is the coexistence of dogmatic and conflicting cultures whose work is more fundamentally guided by the culture’s survival rather than by the growth of knowledge.

To conclude, Longino argues that the influence of ingrained beliefs can only be neutralised by the confrontation of disparate viewpoints. Even though Longino is concerned with ideology and social values, her analysis can be easily extended to any belief that is more or less unreflectively accepted. This is precisely the principle underlying the *argument from sociality*. Having different sets of culture and beliefs participating in the knowledge production process is an effective way of producing epistemically sound novel claims to knowledge.⁵²

⁵² The fact that the SEE does not make explicit mention to other social, cultural, and feminist studies of science does not deny the effect of other cultures or values in science. It implicitly acknowledges this influence in the establishment and stability of dominant scientific cultures and in the determination of scientists’ priors. To be sure, a more ample reading of the sociality argument can account for the influence of these social values. The influence of well-established beliefs, whatever their origin might be, can and should be counteracted by the exercise of effective criticism within and between scientific communities.

SCIENCE AND SOCIETY

The recognition of the epistemic role of the social dimension of knowledge production brought to the fore the importance of the study of the organisation of science and that of the broader interaction between science and society. Not only is the collective aspect of knowledge production relevant to overcome the epistemic difficulties intrinsic to knowledge production, but it is also relevant to overcome the detrimental effect of an unequal distribution of political power in science and society. Kitcher's *Science, Truth and Democracy* (2001) is very illustrative in this respect.

Here, Kitcher addresses the issue of selecting relevant and benign problem-situations to society. The view that takes the attainment of a single and unified account of nature as *the* goal of the individual scientists is now discarded. Scientists aim instead at *significant* truths where significance is determined by scientists' evolving interests, epistemic and non-epistemic alike. Scientists' interests are constitutive of the process of knowledge production in that they determine the selection of research projects and the classificatory concepts and systems by reference to which scientists pursue their inquiries. Instead of providing a single coherent and complete account of nature, or reaching consensus on a unifying theory, science now provides a 'patchwork of locally unified pieces' based on specific and varied classification schemes (2001, p. 72). Nonetheless, Kitcher still allows for a consistent local account of nature grounded on 'relevance relations' established between the complex of (partial and local) truths and the topic of interest (op. cit., p. 75).

The pursuit of scientists' interests is no longer perceived as inoffensive and conducive to the collective good. Insofar as the pursuit of individual cherished values may have undesirable social consequences they ought to be constrained. This is the case when the pursuit of a line of research may interfere with the freedom of others in the exercise of their abilities to define and choose their own projects and aspirations. But this is not an argument for the elimination of certain areas from scientific inquiry. It is rather an argument that points to the responsibility of scientists in weighing the consequences of their research and refraining from those areas which may have a negative impact on others (op. cit., p. 105). This means that the relation between the individual and the collective is reversed. Some individual sacrifice is now required for the attainment of the collective good or at least the avoidance of some collective bad. From this it follows that some benign individual motives are after all important. Though, they are not required for the attainment of the epistemic goals of science. They are instead required for the protection of important values to society. The idea that science is an unqualified good or that knowledge is always beneficial are now at stake. Society can and should refuse the pursuit of lines of inquiry that may jeopardise important collective values. Hence, from the study of the best distribution of cognitive effort, the ultimate goal of social epistemology became the study of the social organisation of science that best promotes "our collective values in the most encompassing sense" (op. cit., p. 111). To this end, Kitcher puts forward the notion of a "Well-Ordered Science" as a conceptual tool that may guide policy-makers and scientists as to how best achieve those goals (op. cit., p. 118).

In a well-ordered science the choice of research projects conforms to the collective good which is defined by the democratic deliberation of an ideal scientific community. The ideal deliberation is undertaken through a process of collective decision-making in which the interests of the various groups of society, including those of future generations and those of less advantageous positions, are taken into account. After discussion, the representatives of society's interests reach consensus on the goals that should be promoted by research. In the possession of the collective goals determined by ideal deliberation and additional information provided by relevant experts, the institutions of science should then allocate resources to the research projects that best achieve the social objectives (op. cit., p. 121). Having set the research agenda, the well-ordered science then promotes the pursuit of the socially desirable projects and their translation into practical consequences.

From an individualistic social epistemology that took individual diversity as pivotal to the accomplishment of the epistemic goals of science, Kitcher's epistemology became more 'holistic' in that it takes as point of departure the social goals defined by a democratic society. The concern now is not the possibility of achieving the epistemic goals of science, but achieving a socially adequate distribution of research projects. The problem is that societies are not perfect democracies and therefore there may be deviations from the ideal distribution of resources and, as a result, moral and social damage may be inflicted on least protected citizens. Nonetheless, the 'well-ordered science' can be a standard by reference to which scientists may guide their practices. In sum:

[T]he narrow function of the sciences is to generate significant truths, where the criteria of significance are those in force within the community of inquirers. The broad function is to promote the democratic practice of science, as conceived in the ideal of well-ordered science (so that, for example, the criteria of significance actually in force represent the outcome of an ideal deliberation among ideal agents). From this broad function flows the responsibility to attempt to lessen the gap between actual practice and the ideal. Beyond that, at an even wider level, the scientist, as citizen, also has the obligation to do whatever is possible to realize more fully democratic ideals in her society (op. cit., p. 195).

From this it follows that the pursuit of a 'well-ordered science' ultimately depends on scientists' judgment. Scientists should be fully conscientious of the degree to which their projects conform to the ideal and contribute to the consciousness of others, including the general public and, in particular, the members of the groups who allegedly may suffer the negative consequences of science. Moreover, scientists should be willing to voluntarily exit projects that deviate from the ideal and replace them by projects that more closely approximate the well-ordered science.

It is perplexing, however, that Kitcher, who explicitly recognises the existence of inequality in society and the consequent selection of projects that accord with the preferences of the privileged, and scientists' tendency to neglect the negative impact of their research projects, leaves the implementation of well-ordered science to the conscience of individual scientists. That is, despite arguments to the opposite effect, Kitcher endorses the view that scientists do know what the ideal of science is and

that they can conform to that ideal. That is, they can overcome their prejudices and privileges and thereby attenuate the negative impact of science on society. But as I have been arguing, there are social, psychological and cognitive barriers to the confrontation of divergent values and interests which the individual scientists are unable to overcome by themselves.⁵³ The removal of such barriers requires special democratic scientific institutions.

3.5 THE SOCIAL EPISTEMIC ARGUMENT FOR CRITICISM

The social epistemological accounts reviewed in the previous section point to the role of the institutions of science in promoting the pursuit of diverse and socially desirable problem-situations and the generation of epistemically sound problem-solutions. However, scientific practice does not conform so neatly to this ideal. On the one hand, diversity is actively and passively pre-empted by the scarcity of resources which concentrate on a far more limited domain of research. On the other hand, the activity carried out in the various fields of research often occurs in isolation. Dialogue amongst researchers from different traditions is a rather difficult endeavour. Not rarely do scientists perceive the work carried out in other fields of research as irrelevant to their own.

Longino noted that dialogue amongst different traditions requires the acceptance of common standards by reference to which the critical discussion is carried out. Failure in reaching agreement on the standards of discussion jeopardises the critical exercise. This difficulty is most explicit in Solomon's epistemology in which dissent is caused by the existence of divergent standards and consensus is explained by the support the different standards provide. But Solomon does not supply a solution to this problem. She simply recommends tolerance towards dissent whenever disputes are empirically sustained and greater caution when analysing epistemic states of consensus. But tolerance and caution do not promote dialogue. They just warrant the co-existence of different and divergent proposals.

SEE explicitly rejects the invocation of cultural differences as an argument against dialogue and critical discussion. To demonstrate the possibility of dialogue among scientists educated in different cultures, it is important to recuperate the conceptions of problem-situation and problem-solution adopted and already extensively illustrated in the present work. As we have seen, the scientific endeavour consists of forging solutions to emergent problems of scientific practice which result from disappointed expectations in the light of some item of scientific culture. The scientific solutions in turn represent successful answers in accommodating the frustrations that gave rise to the problem-situation. On this view, then, *criticism is a challenge to the solving power of some problem-solution that is grounded on a conflict with legitimate expectations from the purview of some other part of scientific culture.*

Given that scientific cultures vary among scientists within a community, and more so among scientists belonging to different communities, the solutions put

forward by different scientists may be different and this difference may be wider the wider the distance between cultures. At the heart of the difficulty at implementing effective and transformative criticism is the fact that the wider the distance between cultures the more scientists will perceive the conflict as non-significant. Not only will they neglect conflicting results, but they will also feel unchallenged by critical charges. But if scientific practice aims at the solution of problems that result from disappointed expectations in the light of established knowledge, the divergence between cultures should not constitute an impediment. This divergence should encourage critical scrutiny, instead.

As the philosophers of science have increasingly recognised, the confrontation of divergent cultures allows the identification and revision of most ingrained beliefs which holds a high potential for the growth of knowledge. Of course, the confrontation of disparate views is difficult. The cognitive, psychological and social costs of submitting to criticism established beliefs create a natural resistance to such endeavour. Scientists tend to take up problems for which they have the expectation that a solution might be forthcoming with the available resources. And a conflict between distant results may be perceived as insurmountable and as result it may be ignored.

However difficult, the exercise of criticism should be encouraged. Criticism does not require the existence of a common ground. Therefore, previous agreement on the 'background knowledge', 'paradigm', 'standards of practice', 'culture' need not be a pre-condition for criticism. What is crucial for criticism is the recognition of a conflict between items of scientific culture. To put it bluntly, whenever scientists' expectations get frustrated there is room for critical discussion. Frustration is epistemically significant. It means that some item of scientific culture is in conflict with legitimate expectations from the viewpoint of another set of items of scientific culture, or it means that a problem-situation has been identified. From this it does not follow that there is no need for shared standards. The existence of a conflict already presupposes a common denominator which renders the conflict cognitively salient. For disjoint items of scientific culture cannot meaningfully conflict with one another.

The exercise of criticism requires that critical challenges are adequate and relevant to the problem-situation and that the criticised view perceives them in such a way. The natural question to ask than is: What is an adequate and relevant criticism? From what I have been arguing it is not difficult to see that an *adequate criticism*, a criticism that ought to be perceived as relevant by those who hold the criticised view, *is one that identifies a conflicting item of scientific culture with a proposed problem-solution.* Scientific practice consists in the forging of solutions to problems that emerge from disappointed expectations where the fruitfulness of the problem-situations depends on the gap of the inconsistency at hand.

To conclude, the view that criticism is a problem-generating practice that results from disappointed expectations from the standpoint of some part of scientific culture solves a major difficulty present in both philosophical literature and scientific disputes. Having consensually shared standards of practice is not a condition for the exercise of criticism. Conflict is epistemically significant by itself. It deserves attention whether it emerges from different cultures, paradigms, research programs, theories, or

⁵³ For a more in-depth criticism along the same lines see Longino (2002b).

experimental evidence. As long as a conflict is identified there is room for criticism. Given the cognitive, psychological and social constraints to the critical debate, important implications derive to the social organisation of science, which should take into account the factors that may obstruct it.

3.6 THE SOCIAL EPISTEMOLOGY OF EXPERIMENT

I am now in conditions of completing the argument from sociality that accounts for the processes whereby scientists produce and establish the results of experiment. I have argued that the epistemic value of the social dimension of knowledge production depends on the heterogeneity of the community of scientists participating in the processes of knowledge production and validation. The heterogeneity of cultures and beliefs allow effacing the constraining effects of, as Popper put it, “those prejudices which are an unavoidable consequence of ... [the scientist’s] peculiar mental history” and “that strange blindness concerning the inherent possibilities of our own results”. However, the mere heterogeneity and large numbers of items of scientific culture do not do the trick. They may, in fact, paralyse the process of knowledge production by giving rise to too many unanswerable questions given the state of the art at a particular point in time. However, the scientists themselves can easily identify the items of culture that can be brought to the process of knowledge production. Their relevance depends on their potential contribution to the resolution of the problem-situation at hand or to the appraisal of the problem-solution under scrutiny. This relevance can only be ascertained by the scientists involved in the process of knowledge production or by those who want to engage in it.

The natural resistance of scientific communities to the intromission of other cultures requires social institutions that promote dialogue among various cultures. These institutions must demand from the newcomers that they be able to justify the relevance of their proposals to the problem-situation at hand. The autochthones in turn must respond to the charges thereby justified. In sum, the potential of the social dimension of knowledge production to generate epistemically sound novel claims to knowledge depends on an adequate organisation of scientific production that encourages the articulation of items from various cultures and openness to engage in problem-situations prompted by the research of others. The impact of science on society, in addition, requires the creation of special forums for divulging the results of science to the general public and, specially, to the parties that may be negatively affect by it. In synthesis, the main results of *The Social Epistemology of Experiment* are:

1. The greater the number and the wider the heterogeneity of the resources (material, conceptual and social) involved in knowledge production, the higher the epistemic status of the three-way coherences and the experimental results supported by them.
2. The variety of the items of scientific culture ensures that relevant questions are likely to be asked and answered during the process of knowledge production.

3. The problem-solutions arrived at are epistemically sound when they are the result of practices that have explored relevant courses of action to finding a solution to the problem-situations at hand.
4. The adequate number and level of heterogeneity of the resources are to be determined by the scientists themselves.
5. The effectiveness of this discussion depends on the organisation of science to promote dialogue between different cultures.
6. The science institutions must also create special forums to divulge the results of science to the general public so that the beneficial effects to society can be extended and the negative impact prevented.

3.7 CONCLUSION

The *argument from sociality* accounts for the way by which experimenters produce and establish knowledge. It asserts that the collective aspect of knowledge production bestows added epistemic value to the three-way coherences by bringing to the production process a large number of items of scientific culture and beliefs. It reinforces the epistemic value of the experimental results insofar as these gain autonomy from the particular set of items of culture and beliefs that brought about the problem-situations.

However in different guises, and to different degrees, the social epistemologies all acknowledge the detrimental effect of scientists’ commitments in knowledge production. And they suggest that the solutions to the problems require social arrangements that promote diversity and effective criticism. These accounts are inspired by the epistemic problems of theory-production and theory-testing. They all take empirical knowledge as given and unproblematic. Therefore, the detrimental effect of scientists’ commitments pertains to the production of epistemically weak theories and to their unjustified protection, or by an inadequate adjudication among alternative theories.

SEE, in turn, focuses, on the one hand, on the possibility of conceptualising fruitful problem-situations and, on the other hand, the production of epistemically sound problem-solutions. The first kind of issues does not find echo in the philosophy of theory. The second issue has a more direct theoretical counterpart. Whereas in theoretical practice, scientists’ priors may lead to the inclusion of *ad hoc* hypotheses and immunising stratagems to re-establish a relation of coherence between a theoretical hypothesis and an empirical result. In experimental science the problem takes the form of overlooking the difficulties that arise when building support for a favourite phenomenal model and the possibility of eliminating conflicting results by way of material or conceptual manipulations.

The focus on the processes by which experimenters produce and establish knowledge, and the resulting reconstruction of experiment as an endeavour that produces three-way coherences, bring to the fore the role of the social process of evaluation in giving rise to fruitful problem-situations. Not only does the collective process of knowledge production contribute to the generation of better problem-solutions, but it also gives rise

to the generation of fruitful problem-situations that emerge from the confrontation of different cultures. In turn, the analysis of the social epistemologies of theory highlights the importance of the institutions of science in promoting both diversity and criticism, which the social epistemology of experiment gladly welcomes.

CHAPTER 4

THE SOCIAL EPISTEMOLOGY OF EXPERIMENT

4.1 INTRODUCTION

The analysis of experiment brought to the fore three arguments that justify belief in experimental claims to knowledge. The first argument is *the argument from coherence* that asserts that the mutual support achieved between the components of the experimental system assures scientists that sense has been made of the object of scrutiny. Insofar as scientific culture and scientists' prior beliefs may limit the array of admissible problem-situations and the adequacy of problem-solutions, the attainment of the three-way coherence does not suffice to justify the epistemic value of experimental results.

Two additional arguments were then identified that reinforce the epistemic value of coherence. The *argument from materiality* asserts that the engagement of the material world in knowledge production frustrates scientists' expectations and thereby promotes the revision of scientific culture and prior beliefs. The *argument from sociality* asserts that the social dimension of knowledge production by bringing to the production process a vast number of items of scientific culture and beliefs promotes the examination of them both. The three arguments in conjunction sustained the conclusion that the greater the number and the greater the heterogeneity of the items of scientific culture (material, conceptual and social), the higher the epistemic status of the coherent results because they are the result of practices that have explored relevant courses of action to solving the problem-situation at hand.

This chapter introduces four tests intended to appraise the epistemic value of the three-way coherence and by so doing completes the Social Epistemology of Experiment. Section 2 presents the four epistemic tests of experiment: 1) the materiality test, 2) the stringency test, 3) the social robustness test, and 4) the technological test. Section 3 addresses the technological application of the results of experiment. Section 4 highlights the import of SEE and section 5 illustrates it by applying the four tests to the Quark research program. Section 6 shows how the relevance of the traditional problems of experiment such as incommensurability, the underdetermination thesis, and the experimenters' regress, is significantly lessened. Section 7 compares SEE with the error-eliminative account of experiment. The conclusion follows in the last section of the chapter.

4.2 THE EPISTEMIC TESTS OF EXPERIMENTS

The epistemic tests of experiments are designed to appraise the epistemic value of the three-way coherence and, in particular, the contribution of the two main epistemic factors: the participation of the material world in knowledge production and the critical nature of the social process of knowledge production.

THE MATERIALITY TEST

The *materiality test* assesses the *material quality* of the experimental system and the *agency of the material world* in the process of knowledge production. It examines the composition of the experimental system and the actual participation of the material world in knowledge production, that is, the actual possibility given to the material world to frustrate scientists' expectations during that process. Of course, the material quality and the material agency of a given experimental system are interrelated. A high level of material agency is associated with a high level of material quality. But as will be shown, an experiment with a high level of material quality may nonetheless possess a low level of material agency. This is what explains the closer scrutiny of materiality.

Three units of analysis are relevant to assess the composition of the experimental system: the *object of scrutiny*, the *experimental practices* and the *outcome of practice*. In regard to the object of scrutiny, the materiality test evaluates whether the phenomenon produced in the laboratory is the subject matter of research or it is instead a surrogate for it. In the latter case many situations are possible. The object of interest may be studied by the examination of a better-known phenomenon which works as a material model of the object of interest, or it might be investigated by having recourse to a non-material model of the phenomenon.⁵⁴ The evaluation of the materiality of experimental practice also requires the examination of the relative contribution of the various procedures that brought about the experimental result. These include material manipulation, conceptual analysis, and data processing techniques (cf. Franklin's list of strategies of experiment, section 2.5). Finally, the outcome of practice might also possess different degrees of materiality. The possibilities here range from the actual production of the phenomenon of interest to its estimation through conceptual manoeuvring via modelling and statistical techniques. The more the experimental system resorts to surrogates rather than the object of scrutiny, and to conceptual rather than material tools, the lower the epistemic value of the experimental results because the lower will be the participation of the material world. The assessment of the actual participation of the material world in knowledge production requires evaluating the extent to which the

experimental result is the outcome of the agency of the 'material world' rather than the actions of the scientists.

Mary Morgan's (2002, 2003, 2005) taxonomy of experiments lends itself as a clear illustration of the materiality test as proposed here. Indeed, this taxonomy seems to discriminate different categories of experiments according to the criteria of material quality and material agency, as defined above. Morgan identifies four kinds of experiments, in ascending order of materiality: '*mathematical model experiments*', '*virtually experiments*', '*virtual experiments*' and '*ideal laboratory experiments*'. The 'ideal laboratory experiment' is the scientific device with the highest level of materiality whereas the 'mathematical model experiment' is at the lower end of a continuum of 'vicarious' experiments, which may be classified either as 'virtually experiments' or 'virtual experiments'.

The 'ideal laboratory experiment' corresponds to the stylised view of experiment which consists in the creation and control of a material object with the goal of observing it under the favourable conditions of the laboratory (cf. section 2.2). Morgan locates the material component of experiments at three distinct levels: the composition, the mode of control, and the method of demonstration of the experimental system. In 'ideal laboratory experiments' the inputs, the experimental process, and the outputs are material. Control is achieved by materially interfering with the object of scrutiny and by materially shielding it from the surrounding environment. And the results obtained derive from the direct observation of the object under scrutiny. In contrast, 'mathematical model experiments' consist in the generation of logical deductions from well-defined premises. In these experiments, the inputs, the mode of intervention, and the outputs are mathematical; control is exercised via simplifying assumptions (e.g. *ceteris paribus* conditions, assumptions of independence between variables, etc.) and the results are the outcome of deductive or logical manipulation of a mathematical system. 'Vicarious experiments' are hybrid devices that combine both 'model' and 'experimental' elements in the generation of scientific results. Their distinguishing feature is that while they rely on a process of demonstration closer to experimental intervention in the sense that they involve the analysis of an external intervention upon the object of scrutiny, they use mathematical resources as surrogates for material processes. Vicarious experiments are 'virtually experiments' if the intervention is made on semi-material objects; but if the object of scrutiny is nonmaterial they are instead 'virtual experiments'. According to Morgan, semi-material objects are models of the object of interest that retain a high degree of structural 'verisimilitude' with the object of reference, and non-material objects are entities with a more idealised and simplified abstract structure. For this reason, 'virtually experiments' are closer to 'ideal laboratory experiments' while 'virtual experiments' approximate 'model experiments'. Thus, on Morgan's account the ideal laboratory experiment possesses a higher level of material quality and material agency in comparison with other experiments.

In sum, the higher the level of materiality, the higher the epistemic value of the experimental result because the knowledge production process conveys a high chance of producing new knowledge about the object of scrutiny. This possibility stems from the ability of the material world to frustrate scientists' expectations.

⁵⁴ The model literature has already addressed this issue quite extensively. Classic references include Mary Hesse (1966); Rom Harré (1970); Max Black (1962); Peter Achinstein (1968). More recent contributions can be found in Mary Morgan and Margaret Morrison (1999). See also Uskali Mäki (2001, 2005). This issue is addressed in more detailed in Ch. 8.

THE STRINGENCY TEST

The *stringency test* assesses the rigidity of the experimental systems, that is, it evaluates the extent to which the results of experiments may be rendered compatible with an item of scientific culture, or an expectation, by material or conceptual manipulation of some part of the experimental system. A scientific result is a stringent result if it is produced by a rigid experimental system that is not easily amenable to material and conceptual manipulation. The stringency of an experimental result increases as other material apparatuses, grounded on well-established theories, reproduce it. This is the case because the robustness of an experimental result to changes in the experimental system increases its autonomy relative to the material and the conceptual means of production.

Whereas the materiality test assesses the kind and the level of participation of the material world in knowledge production, the stringency test assesses the epistemic value of this participation. The epistemic value of the participation of the material world is higher the stronger the degree of rigidity of the experimental system, or the lower its level of plasticity.

THE SOCIAL ROBUSTNESS TEST

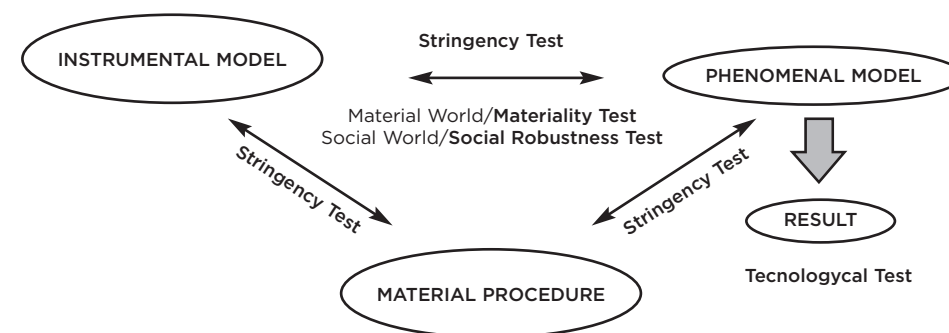
The *social robustness test* evaluates the social process that gave rise to the experimental result. It calls for judgement about the critical nature of the context wherein the experimental claim to knowledge was produced and established. The goal is to assess the extent to which the experimental process of knowledge production answered relevant questions to the problem-situation at hand or, to put it in another way, the extent to which the experimental result is supported by varied and heterogeneous items of scientific culture. On this view, a socially robust scientific result is one that has undergone an intense process of criticism and that has successfully addressed critical charges levelled against it. The term ‘robustness’ is here intended to emphasise the epistemic import of the social process of evaluation to the ‘reproducibility’ of the experimental result. The upshot is that the collective acceptance of an experimental result as a valid result assures that it is highly replicable and robust. It provides further assurance that other scientists can or believe they can reproduce it irrespective of their different cultures and beliefs.

THE TECHNOLOGICAL TEST

The *technological test* assesses the degree of applicability of the products of scientific experimentation to new problem-situations. Given that the degree of applicability of the products of science is often associated with technological success, which is taken as the hallmark of science, this test is labelled here the technological test. It therefore evaluates the capacity of the experimental process of knowledge production to generate novel items of scientific culture. The applicability of the products of science is epistemically significant insofar as it improves their technical and social robustness. It does that by extending the autonomy of the experimental results from the cultural and

the social context in which they were produced and established. The technological test is also a measure of the social utility of the products of science, in science and elsewhere. Because the applicability of the experimental outcomes was not addressed in the previous chapters, it will be the topic of the next section. Before examining the technological success of the products of experiment, the function of the experimental tests is represented in the scheme below.

FIG. 4.1. The three-way experimental coherence and the epistemic tests of experiments



From figure 4.1 it is clear that the four tests are interdependent and that they reinforce each other. The stringency test assesses the extent to which the conceptual models and the material procedure fit together. Of course, the stringency of this fit improves if the experimental system incorporates a high level of materiality and if a critical scientific community brought it about. The technological test, however, offers only an indirect test to past practice. By tracing the subsequent path of experimental results it nonetheless provides a test to the fruitfulness of previous problem-situations.

4.3 EXPERIMENT AND TECHNOLOGICAL SUCCESS

The success of scientific experimentation is often associated with the technological innovations it brings about and that have substantially altered human life. This is not surprising given that scientific experimentation is to some extent a technological achievement in that it attempts to produce reliable equipments tailored to intervene in the material world in predictable and stable ways.⁵⁵ However, very little has been said here about what happens to the products of experiment once produced. Thus far the concern has been with the production of an experimental result and its acceptance by the wider community. This section is intended to fill this gap. It analyses what Thomas

⁵⁵ The machine analogy has been extensively used by Nancy Cartwright and the conception of experimenting as engineering has been explored by the philosophers of experiment (cf. Cartwright 1999, and Janich 1978, see Appendix 2).

Nickles calls “second-order work on problem solutions and other results already gained [rather] than the work of constructing those results in the first place”, the “continuation of construction” or “reconstruction” (1989, p. 310).

[C]onstruction concerns the original manufacture of knowledge claims (usually in the form of published articles), and their immediate interpretation and acceptance (or rejection) by the relevant scientific community. Reconstruction concerns what happens to these results later – how they are streamlined, articulated as nodes of well-understood frameworks, and, in some cases, ‘black boxed’ so that they can be routinely used without full understanding. (ibid)

At this juncture, a distinction should be made among the three final products of experiment: the material procedure, the instrumental model and the phenomenal model. As suggested in chapter 2, the final versions of material procedures and instrumental models that underlie the production of experimental results yield devices, instruments and machines that generate stable outcomes when operated in a regular manner under similar conditions. The two-way coherence achieved between the material apparatus and the instrumental model defines its appropriate operation for specific purposes. However, the use of material apparatuses for other purposes may lead to further work on the material apparatuses and their respective instrumental models. If it does not, the material apparatuses become ‘black boxes’, i.e. routinely used apparatuses the full understanding of which are no longer necessary.

In the science studies the well-established items of scientific practice have been labelled ‘black boxes’ (e.g. Pinch 1986, p. 212-4; Latour 1987, p. 2-3, p.131) to express the autonomy of scientific results from their context of production. When a scientific product becomes a black box it becomes a reliable object that can be routinely applied without its users being conscious of the process (technical and social) that brought it about. Thus, from the viewpoint of the users, and when referring to material apparatuses, a black box is a complex object of which it is only necessary to know its input and output.⁵⁶ The same holds for the adoption of ‘phenomenal models’ for other scientific purposes. When phenomenal models are adopted for the solution of new problem-situations, either in experimental practice or in theoretical development, they become factual statements detached from their production processes. They become facts whose status is no longer in question. But scientific facts, too, presuppose the interactively stabilised material procedures and conceptual models that sustain them.

In theoretical construction, for instance, the ‘fact’ becomes part of an argument with a well-defined structure, an internal ‘logic’, which may be independent from

the procedures that led to its discovery.⁵⁷ That is, the fact becomes empirical evidence for a theoretical argument of a deductive or inductive kind.⁵⁸ However, experimental facts can be used for theoretical arguments only after considerable practice has been undertaken and both experimenters and theoreticians have become full experts in the experiments that generated the results. When these arguments are finally accepted they have sound epistemic value conferred to by the force of well-established facts and the coherent account within extant scientific culture. The ‘fact’ has now become an item of scientific culture that conveys epistemic value to the novel products of science.

In sum, subsequent scientific practice consolidates previously established results by adopting them in new problem-solving activities. But it should be noted that insofar as these results are brought to bear on new problems, their understanding evolves and, consequently, they may be modified, refined, improved, as well as rejected by subsequent practice. In the process, weakly defined problems, weakly established techniques, and conjectural solutions may be reinforced and substituted for better-structured problems, techniques, and empirically established theories.

4.4 ON REPLICABILITY AND ROBUSTNESS

The concepts of replicability and robustness are the traditional epistemic criteria of experiment. *Replicability* is the requirement that an adequate experimental result ought to be reproduced whenever the original material procedures are themselves reproduced. This requirement is meant to ensure that the created phenomenon is not a fortuitous event or an artefact of the experimental procedures. But this does not preclude the fact that a phenomenon may be reproduced in various different ways. The other epistemic concept captures this possibility. *Robustness* expresses the added epistemic value of a result that is reproduced by alternative material procedures and their respective conceptualisations.⁵⁹

Replicability is a requirement of good experimental practice that alludes to the private stage of knowledge production while robustness appeals to its public stage. Experimenters must ensure that their results are reproducible by themselves and by other scientists were they to encounter the same material and conceptual conditions. Yet replication is seldom realised by other scientists (cf. Mulkay and Gilbert 1986, Collins 1985).⁶⁰ On the one hand, the creation of a novel phenomenon involves a long learning process whereby experimenters gather information about the experimental system which is not generally available to other experimenters. Therefore, straightforward replication

⁵⁷ Logic here is not to be confused with deductive reasoning, nor with a set of well-established algorithms nor even with routine heuristic procedures.

⁵⁸ Nickles talks about ‘consequential arguments’ which presents the ‘fact’ as the predictive consequence of a theory and ‘generative arguments’ which take the ‘fact’ as the empirical source for theoretical construction.

⁵⁹ These definitions capture the philosophers’ terminological use of the concepts. Scientists tend to conflate and use them interchangeably.

⁶⁰ That replication rarely occurs in scientific experimentation was not only pointed out by the sociologists of science, but it was also acknowledged by other students of experiment and by the experimenters themselves.

⁵⁶ Hans Radder (1995, p. 85) argues that for a system to become a black box there is no need for scientific consensus about the system’s inner workings. The epistemology of experiment developed here argues instead that in order to become a black box, a scientific result must first be established as a proper solution to a problem, which requires the three-way coherence and therefore consensus between the relation between the material procedure and respective conceptual models. The utilisation of the apparatus for other purposes requires the respect of the instructions of proper operation as described by the instrumental model. However, as long as these instructions are followed, the knowledge of the inner workings of the material apparatus is not required.

may be unreachable to those who were not directly involved in the production process. On the other hand, and more importantly, there is no great incentive to generate a replication except when the experimental result is controversial. Experimenters are more often interested in investigating the effect of purportedly relevant modifications in the experimental set-up. Thus, replicability is better understood as a requirement that asserts that experimental results ought to be replicable, rather than a request to actual replication. Robustness, in turn, presupposes the actual reproduction of scientific results under different material and conceptual conditions, preferably performed by other scientists.⁶¹

Just as the requirement of replicability does not have to be actually realised, there can also be a potential interpretation of robustness. The appraisal of experimental results by scientists other than those involved in the production process enhances the robustness of the results even if this appraisal does not involve the material reproduction of the result. This is the sense in which the social dimension of knowledge production adds and complements the epistemic value of experimental knowledge and which has been overlooked in the studies of experiment.

The analysis of experiment undertaken here supports the criteria of replicability and robustness as adequate for the appraisal of experiment. They are also measures of the effect of scientific culture and scientists' priors in the experimental process of production of knowledge. In this regard, replicability is *grosso modo* a demand for the independence of the results from the personal idiosyncrasies of the scientists while robustness is *grosso modo* a demand for their independence from scientific subcultures. This independence is then the underlying principles of the criteria of replicability and robustness. However, these concepts, as traditionally understood, do not account for all the relevant factors and their relative epistemic import. They are silent about the contribution of material agency to knowledge production and that of the social dimension of knowledge production, which the stringency, the materiality and the social robustness tests assess.

4.5 THE HUNTING OF THE QUARK 4: THE FINAL TEST

The import of the proposed social epistemological framework is now illustrated by way of applying the four epistemic tests of experiment to the quark research program and by comparing its conclusions to those derived by Pickering.

THE MATERIALITY TEST

In the quark research program two equivalent material systems produced contradictory results. The Stanford and the Genoa experiments belonged to the same category of experiments which were built upon the same resources of practice – the Millikan experiment and classical electromagnetic theory.⁶² Both groups of scientists

built specific material apparatuses to detect quarks through direct measurements of electric charges on samples of matter. Even though each set of experiments involved specific material and conceptual practices, the object of scrutiny, the procedures and the outcome of the practice possess comparable levels of materiality. That is, the material quality of these experiments is comparable. Insofar as both sets of experiments attempted to realise direct measurements on samples of matter, the material agency of these experiments is also equivalent. In sum, the score of the materiality test of both sets of experiments is high.

THE STRINGENCY TEST

The Stanford and the Genoa experiments resorted to the same well-established resources of scientific culture. The Millikan experiment had already achieved the status of a well-established item of scientific culture. It was credited to have established the primacy of the charge on the electron and its absolute value. Classic electromagnetic theory was a respected theory that asserted that in a uniform electric field, the electric charge is the only force acting on a sample of matter. The conjunction of the two resources rendered feasible the calculation of the electric charge of matter from the measurement of the mass of matter and the voltage of the electric field.

But the search for quarks required the inspection of larger samples of matter than the Millikan apparatus allowed. The Genoa and the Stanford groups had to build new apparatuses to that end. The necessity of creating a new apparatus for the search of an unknown object rendered both sets of experiments highly plastic. The plasticity of the experimental systems therefore reduced the role of the material world in knowledge production. This is, in fact, what explains the production of the two conflicting results.

The central difficulty of these experiments consisted in determining whether a perfect uniform electric field had been created. Consequently, an inconsistency could always be attributed to the presence of inhomogeneities in the field and, consequently, it could be removed either by material or conceptual manoeuvring. As it was shown, whenever the Genoa group found non-integral charges, the material apparatus was calibrated to conform to classical electromagnetic theory. In particular, when the group found grains that presented non-zero residual charges, the apparatus was adjusted so that the measurements were made consistent with zero charge (cf. section 2.4). The Stanford group, in turn, always succeeded in reproducing fractional charges (cf. section 2.9). Evidence for the existence of isolated quarks could thus be either supported or challenged on the grounds that the results were caused by inhomogeneities in the electric field.

This case study suggests that when the material procedures and the instrumental model are both very plastic resources of practice, the re-establishment of a relation of coherence seems to be highly dependent on the prior conception of the phenomenon under investigation. Both groups of researchers succeeded in replicating the first phenomenal model in subsequent experiments. Even though there is no reason to believe that each experimental group had an *a priori* preference for either phenomenal model, these experiments do suggest that subsequent action was informed by previous results. Insofar as the two sets of experiments relied on

⁶¹ For a somewhat different conceptualisation of these notions see Radder (1995).

⁶² They belonged to the category of stable-matter experiments. Other types of experiment included particle-accelerator experiments and cosmic-ray experiments.

equally plastic experimental systems, the level of stringency was low in both. In sum, both the Stanford and the Genoa experiments exhibit a low level of stringency due to their highly plastic experimental systems.

THE SOCIAL ROBUSTNESS TEST

The social dimension of knowledge production shows up very clearly in Pickering's account of *The Hunting of the Quark*. Even though the problem-situation was restricted to two phenomenal possibilities, which limited the admissible courses of action, and despite the fact that one of these models received stronger support from extant culture, the fact that both courses of action were pursued is epistemically significant. Thus, the obstructing effect of the scientific culture and the scientists' priors at an aggregate level was mild. Moreover, the two groups of experimenters did want their results to be accepted by the wider community. This is particularly clear in the Stanford group. Inasmuch as their results were in conflict with established knowledge they had to make sure that their results could not be attributed to spurious effects. This explains why it took the Stanford group twelve years to claim that they had found evidence for the existence of isolated quarks. The Genoa group, in contrast, published their consistent results with established knowledge two years after the search for quarks began.

As the various episodes of *The Hunting* have shown, the problem-situation of finding isolated quarks involved a wide community of scientists, experimenters and theoreticians alike, with varied scientific subcultures. This means that the results of the Stanford and the Genoa experiments were appraised from a variety of cultural standpoints and scientists' commitments.

Pickering rightly notes that the process of collective appraisal is biased towards the "socially sustained matrix of commitments" in that "scientific communities tend to reject data that conflict with group commitments and, obversely, to adjust their experimental techniques and methods to "tune in" on phenomena consistent with those commitments (1981, p. 236). This was the case in the quark program. This is supported by the relative ease with which the Genoa group publicised their results, the contrasting difficulty faced by the Stanford group, the scepticism of the scientific community towards the Stanford claim to have found evidence for quarks and its acquiescence under the unifying tungsten hypothesis. But Pickering fails to notice the epistemic role of the social dimension of knowledge production. He only acknowledges the role of the 'socially sustained matrix of commitments' in constraining experimenters' practices and thereby preventing the production of results that "suit their fancy" (1981, p. 236).

In chapter 3 it was shown that the criticism the Stanford group encountered gave rise to new problem-situations the solution of which increased the information available to the resolution of the conflict. The fact that only the Stanford experiments generated evidence for the existence of quarks after two decades of practice was epistemically significant. And this significance was brought about by the collective work of various subcommunities of scientists with partially overlapping and partially distinct scientific cultures and beliefs regarding the problem-situation at hand. The

point is that even though individual sets of experiments are dependent on particular subcultures and on scientists' commitments, the confrontation of the various contributions is possible and endows the collective consensus with additional epistemic significance: a significance that results from the autonomy gained from scientific culture and the particular idiosyncrasies of scientists or groups of scientists.

THE TECHNOLOGICAL TEST

Pickering's account does not provide enough information about the impact of 'The Hunting' in science, or elsewhere for that matter. The research program ended despite the absence of unanimity among experimental and theoretical researchers. Twenty years later it appears to be no more than a scientific curiosity. However, there are reasons to believe that the quark program did not expire with its last sigh. The experimental and theoretical practice generated several products which were subsequently adopted in the resolution of new problem-situations: technical knowledge acquired with the development of new material apparatus (e.g. reported in journals and PhD dissertations), the phenomenal knowledge acquired by creating new effects in the lab (e.g. the discovery of the magnetoelectric force by the Genoa group), the ability to materially control unanticipated effects, and so forth. This means that the research program continued to influence the practice of the scientists involved and potentially the practice of others who came to interact with them in subsequent practice. These further applications increased the epistemic value of the products of the research program by improving their autonomy from the original context of knowledge production.

4.6 THE LACK OF A LOGIC OF EXPERIMENT

The absence of a logic intrinsic to experiment, which was at the centre of the sociologists' radical criticisms, can now be put into a new perspective. This can be done by drawing the implications of the main critiques that challenged the traditional role of experiment: the incommensurability of experiments, the underdetermination of experimental systems, and the problem of the experimenters' regress.

ON INCOMMENSURABILITY

The analysis of scientific experimentation undertaken thus far shows how each experimental product is a local achievement contingent to the items of scientific culture and practices that produced it. The dependency of the experimental products on the context of production may, in turn, create obstacles to intersubjective scrutiny.

Three different positions were reviewed in the previous chapters. A middle ground stance (e.g. Pickering's) asserts that there is an indefinite number of equally valid problem-solutions, each of which depends on the actual actions performed while attempting to manipulate the material world. On this view, the impossibility of adjudicating among those solutions does not deny that each captures some aspect

of the engagement with the material world. Nor does it deny the possibility of dialogue. Different communities may understand the work operated by the others. The point is simply that they generate information about different aspects of the world that the specific engagement with the material world revealed.

A stronger stance (e.g. Hacking's) derived bolder implications. Experiments generate outcomes that are not destabilised by subsequent ones. The products of laboratory experiments are superseded by others that answer to different problem-situations. The bulk of scientific endeavour is devoted to the extension of knowledge within the borderlines defined by a well-defined domain of research which does not challenge the resources nor the results of prior practice. Once established, theories and instruments remain valid in their data domain.

Finally, a weaker notion of incommensurability, the one endorsed here, asserts that in spite of the level of incommensurability of different experimental systems, the confrontation of different experimental systems is not impossible. This possibility derives from the emergence of frustrated expectations. The main argument of the present work is in fact that alternative problem-solutions ought to be tried and confronted so as to promote the scrutiny of most ingrained and implicitly held beliefs. The confrontation of results is also a local matter which is to be resolved by scientific practice rather than by abstract argumentation.

ON THE UNDERDETERMINATION THESIS AND THE EXPERIMENTERS' REGRESS

The underdetermination thesis asserts that theory is underdetermined by empirical evidence because an infinite number of theories can be made compatible with any given body of data. A conflict between data and the theoretical hypothesis under test may always be attributed to some auxiliary hypothesis rather than to a failure of the theory. As a result, there is no such a thing as conclusive empirical test. The experimental version of the underdetermination thesis is the "experimenters' regress". It asserts that no experimental result is conclusive because it may be due to an artefact of the experimental procedure which cannot be determined by conducting more experiments which themselves may be prone to artefacts.

The analysis developed in this chapter stresses the significance of the underdetermination problem to plastic experimental systems. Under very plastic circumstances an inconsistency between components of the experimental system may be resolved in an indeterminate number of ways, either materially or conceptually. Therefore it might always be possible to generate data that conform to a preferred phenomenal model. In rigid experimental systems, conceptual and material manoeuvring with the experimental system is more difficult. Therefore, the seriousness of the problem is significantly reduced. However, the analysis developed here also revealed the effectiveness of the social process of evaluation in confronting and discriminating alternative proposals. This is an aspect that the study of the actual processes of knowledge production brought to the fore. The point is that even though a single experiment is not by itself sufficient to provide support for the phenomenal model, the set of experiments carried out by the community of researchers do narrow down the range of admissible interpretations for the experimental result.

4.7 COHERENCE AND ERROR ELIMINATION

The most elaborate and consensual account of experiment to date sees experimentation as an activity that consists in the progressive elimination, measurement and calculation of 'background effects', 'noise', or 'error', until experimenters are confident that the phenomenon created was successfully isolated from disturbing factors (cf. Franklin 1986, 1990, 1999; Galison 1987; Hon 1989, 2003; Mayo 1996; Guala 2005b). Accordingly, as errors are eliminated, corrected and accounted for, belief is reinforced as to the 'reality' of the 'signal' looked for. The process stops when experimenters are convinced they have accounted for all sources of error.

The coherence account provided here is not incompatible with the error-elimination account of experiment. The account it provides focuses on the accomplishments of the experimental process of knowledge production. Rather than focusing on what the experimental result is not (i.e. an artefact of the experimental procedure), the social epistemology focuses on what it is – an intelligible result supported by a coherent resolution between practice, outcome of practice and their respective conceptualisations. Rather than being conceived of as results about which experimenters feel confident that they cannot be attributed to a failure of the material procedure, SEE takes the products of experiment as intelligible results that experimenters think they have understood in the course of experimental practice.

The conception of experiment as an error-elimination endeavour does not illuminate the source of the epistemic value reached at the end point of the experimental process of knowledge production, at least not beyond the confidence that the phenomenon of interest has been created. That scientific experimentation consists of the progressive elimination of error until experimenters feel confident that they have eliminated all spurious effects is not very informative. How these effects are identified and eliminated in the course of practice must also be part of the inquiry. And this involves answering questions such as: What was the participation of the material world in the production of the results? Were the experimental systems stringent systems? Did the process leading to the establishment of the results involve the relevant scientific community? These questions are relevant to an eliminative account of experiment insofar as the constraining effect of scientific culture and the biasing influence of scientists' priors may reduce scientists' ability to identify and eliminate error. That is to say, additional arguments must be provided as to why the error elimination procedures have accounted for relevant errors.

SEE provides such arguments. The reformulation of the central epistemological problem-situation of scientific experimentation brought to the fore the central factors and how these factors contribute to the epistemic worth of the results of experiment. As argued and explained above, the 'materiality' and the 'sociality' of the experimental process of knowledge production give rise to interesting problem-situations and robust problem-solutions which are the basis of the growth of knowledge. Inasmuch as the level of materiality and sociality of the experimental processes might vary substantially, the analysis of scientific experimentation must evaluate the extent to which the material and the social worlds have participated in the knowledge production process.

Even though the material and the social aspects of knowledge production show up in the error-elimination historical accounts of experiment, no implications are derived from them. These accounts are still too close to the traditional depiction of experimentation as a linear sequence of operations where the focus is placed on the end products of experiment – an uncontaminated material apparatus and a well-identified and stable natural phenomenon. SEE focuses on the processes of knowledge production and on their bearing to the epistemic value of experimental results. It is also suitable to the appraisal of the ongoing processes of experimental knowledge and tentative experimental results.

4.8 CONCLUSION

The first part of the present work is concluded. A social epistemological framework for the descriptive and normative analysis of scientific experimentation has been proposed. This framework was built upon an informed understanding of how scientists produce and establish experimental claims to knowledge which was derived from philosophical, historical, sociological, and anthropological studies of experiment.

The analysis took as point of departure Pickering's conception of scientific experimentation as the making of 'three-way coherences', which was found pervasive in other accounts of experiment. Experimental practice was then conceived as the forging of relations of coherence among heterogeneous items of scientific culture until a three-way coherence is achieved among the basic experimental components: the material procedure, the instrumental model and the phenomenal model. Coherence grounds experimenters' beliefs in experimental results. The mutual support of the material procedures, the instrumental model and the phenomenal model help forming an intelligible system that gives experimenters' confidence in their practice and in the outcomes of that practice.

But the three-way coherence does not suffice to warrant belief in experimental results. Scientists often select problems for which they foresee a potential solution. This means that the problem-situations are restricted to a narrow domain and that problem-solving is informed by scientists' attempts at implementing an anticipated solution that may lead to the premature ending of the research whenever expectations are fulfilled. Two additional arguments have been identified that reinforce the epistemic value of the relation of coherence underlying experimental results. The *argument from materiality* highlights the epistemic role of the direct participation of the material world in knowledge production. The *argument from sociality* highlights the epistemic role of the social dimension of knowledge production. The implication of the three epistemic arguments of experiment is that the greater the number and the greater the heterogeneity of the resources (material, conceptual and social) involved in knowledge production, the higher the epistemic status of the relations of coherence established because they are the result of practices that have explored relevant courses of action to the resolution of the problem-situation at hand. The epistemic tests of experiment aim at assessing just that – the epistemic status of the relations of coherence that ground the novel claims to knowledge.

In the second part of the present research, I will apply the four epistemic tests to the method, arguments, social processes, single experiments and series of experiments of economics. This application will then supply a test to the relevance and significance of SEE.

PART TWO

THE SOCIAL EPISTEMOLOGY OF EXPERIMENTAL ECONOMICS

CHAPTER 5

THE MARKET EXPERIMENT AND THE FOUNDATION OF EXPERIMENTAL ECONOMICS

5.1 INTRODUCTION

In the previous chapter a social epistemological framework was proposed for the analysis of the processes and products of scientific experimentation. In this chapter the proposed framework is applied to the stylised view of a microeconomic experiment.⁶³ The aim is to provide a general portrait of the experimental method of economics to set the stage for the more in-depth analysis carried out in subsequent chapters. This will be accomplished in two steps. First, this chapter examines the path that Vernon Smith traced in his attempt to put forward an account that explained and justified experimenting with laboratory microeconomic systems. The underlying presupposition is that the epistemological appraisal of any product of science must take into account its actual production process and the context wherein it took place. Second, the chapter appraises an exemplar of a microeconomic experiment to set the standard by reference to which other types of experiments are to be compared with in the remainder of this thesis.

Vernon Smith is the field's prominent practitioner who significantly contributed to the construction and establishment of a particular branch of the experimental field of economics – the market experiments. The study of Smith's trajectory is particularly illustrative in showing the difficulties scientists face when exploring a new area of research. Its establishment involves a long process of knowledge production until the significance of the novel products stabilise and arguments can finally be presented to justify them. This is not to say that Smith was and has been a lonely explorer. Smith's journey was certainly not a solitary one. As will become

⁶³ Two grand strands of experimental work in economics can be identified. One field of research, mostly carried out in the United States, is strongly influenced by an 'outcome orientation' that focuses on the study of aggregate outcomes at the market level. The other research field is mostly carried out in Europe. It follows the tradition of Behavioural Economics emphasising instead the study of the cognitive processes of individual decision-making and behaviour (cf. Friedman and Sunder 1994, p. 127). However, this demarcation has become blurred in recent years as experimenters of a more 'market oriented' bent have come to explore problem-situations of the 'behavioural' kind. Symptomatic of this is the increasing publication of the latter genre of experiments in 'mainstream' journals such as *Econometrica* and *The American Economic Review*. The growth of the field of Behavioural Economics in recent years also testifies for the wider recognition of the relevance of the behavioural studies to economics. The stylised view of a microeconomic experiment belongs to the first field of research which is the focus of the present chapter.

apparent from the historical account, Smith's contribution was as informed by the work of those who preceded him, as by the work of those who collaborated with him, and perhaps most importantly, by the resistance of those who created obstacles on his way to establishing the experimental field in economics.

The chapter is organised as follows. Section 2 presents the stylised view of a microeconomic experiment. The subsequent sections follow Smith's trajectory as he attempted to supply a coherent account of the experimental method of economics. Section 3 starts by presenting Smith's first market experiment and section 4 places it into the broader context of the experimental practice that was emerging at the time. Section 5 returns to the analysis of the market experiment to present the set of arguments Smith adduced in support of both the experimental method and the original experimental results. Section 6 appraises Smith's market experiment by applying the four epistemic tests of experiments. Section 7 shows how the conception of the first market experiment evolved in subsequent practice until it became an established item of the scientific culture of experimental economics. The necessity of justifying both method and results required the formulation of compelling arguments that could persuade the community of economists about the adequacy of the experimental method for economics. Section 8 shows that in order to do that Smith had also recourse to items from other scientific cultures. Section 9 concludes.

5.2 THE STYLISED VIEW OF MICROECONOMIC EXPERIMENTS

This section provides a stylised view of microeconomic experiments. It defines the microeconomic system and describes how economists create and control microeconomic systems in the laboratory. The goal is to present the main concepts, procedures, and purposes of economic experiments to set the stage for the analysis carried out in the second part of this thesis.

The stylised view of a microeconomic experiment is built upon the seminal article 'Microeconomic Systems as an Experimental Science' which Smith published in 1982 in *The American Economic Review*. This view articulates the practice and the conceptual understanding of the practice that accumulated in the field since the late 1950's. The time span of about two decades was necessary for experimenters to become experts in the experimental method so as to be able to provide a coherent account of the new field of research. This is how Smith describes the situation in 1982:

The experimental literature contains only a few attempts to articulate a "theory" of laboratory experiments in economics (Charles Plott, 1979; Louis Wilde, 1980; my articles, 1976a, pp. 43-44, 46-47; 1976b; 1980). It is appropriate for this effort to have been modest, since it has been more important for experimentalists to present a rich variety of examples of their work than abstract explanations of why one might perform experiments... This seems to be the time and place to attempt a more complete description of the methodology and function of experiments in microeconomics. (p. 923)

This article therefore represents Smith's effort to come up with an account of laboratory experiments that could provide a more comprehensive depiction of the experimental field of economics. Experimental economics is here depicted as the field of research concerned with the creation of controlled microeconomic systems in the laboratory for the purpose of measurement of variables relevant to the study of human behaviour and microeconomic systems.

DEFINING THE MICROECONOMIC SYSTEM

The microeconomic system is composed of two elements: the environment and the institution. The *environment* specifies the number and the attributes of the economic agents operating in the microeconomy. These attributes characterise the economic agents' preferences, information, technology, and commodity or wealth endowments. The *institution* consists in the set of rules under which the economic agents communicate, exchange, or transform commodities for the purpose of modifying the initial endowments in accordance with their preferences. More specifically, the institution comprises: the *language* that defines the set of admissible messages that individuals may use in communication; the *adjustment process rules* that govern the sequence and exchange of those messages; the *cost imputation rules* that assign the costs to be applied to the messages; and, finally, the *allocation rules* that compute the outcome of the exchange process (Smith, 1982, p. 924-5). To give an example, the first-price sealed-bid auction is an institution that governs the sale of a single item. It begins with a request for bids after which each buyer submits one bid price. When all bids have been received the allocation rule determines that the item should be awarded to the highest bidder who buys the item at his bid price.

The microeconomic system's performance depends on both the individual agents, who choose messages in accordance with their characteristics, and on the institution, which determines how those messages translate into final outcomes. Given that the performance of the laboratory microeconomy is determined by the conjoint effect of the rules specified by the institution and the decisions of the agents operating under those rules, the microeconomic system is examined by investigating the effect of varying one of the system's variables at a time while keeping the rest constant.

Of course, this is a very general and abstract conceptualisation of a microeconomic experiment. But it is a useful framework that organises the various kinds of microeconomic experiments. In this respect, three grand categories of experiments may be found in economics (Davis and Holt 1993; Roth 1995a). First, there are *market experiments* where subjects perform the roles of buyers or sellers (or producers and consumers) in the exchange (or production) of a particular good under a well-defined set of market rules. Second, there are *game theory experiments* where subjects engage in problems of cooperative or strategic interaction. The major difference between these two types of experiments is that while the market experiments aim at studying the rules of exchange and communication, which are analysed by reference to performance criteria such as traders' surplus, the game experiments aim at studying individuals' behaviour which is analysed by reference to selected rationality criteria, such as Nash equilibrium. Or to put it in another way, while the former focuses on the

study of microeconomic institutions, the latter focuses on the study of microeconomic environments. But in both kinds of experiments individual outcomes are determined by the joint effect of the set of rules and subjects' interactions with other individuals.⁶⁴

Finally, there are *individual behaviour and decision-making experiments* where individuals take decisions in contexts of uncertainty or risk. Here the goal is to study individuals' preferences or appraise individuals' decisions by reference to some pre-defined standards (e.g. consistency). In these experiments, the final outcome for the individual depends only on the decisions she made and the states of the world that emerge after the decisions were taken (this type of experiment can also be described as game against nature).⁶⁵ Examples of the first kind of experiments are provided in this chapter, and examples of the other two are supplied in chapters 7 and 9, respectively.

CREATING AND CONTROLLING A MICROECONOMIC SYSTEM

In chapter two above, two main ideas were included in the definition of an experiment: the *creation* and *control* of a phenomenon (section 2.2). In economics, experimenters seek to create and control microeconomic systems in the laboratory. This is achieved by designing microeconomic institutions and by making sure that these are meaningful to the subjects in the intended way. According to Smith, this is achieved “by using a reward structure and a property right system to induce prescribed monetary value on ‘abstract’ outcomes” (1982, p. 931) and by ensuring that these satisfy the “sufficient conditions for a valid controlled microeconomic experiment” (p. 930): nonsatiation, saliency, dominance and privacy.⁶⁶

Nonsatiation dictates that the amount of the reward earned by participating in the experiment must be important to experimental subjects: “Given a *costless* choice between two alternatives, identical... except that the first yields more of a reward medium... than the second, the first will always be chosen... over the second, by an *autonomous* individual” (p. 931, emphasis in original).⁶⁷ *Saliency* imposes that the amount of that reward must be related to the outcomes of the individual actions in an appropriate way: “Individuals are guaranteed the right to claim a reward which is increasing (decreasing) in the goods (bads) outcomes” (p. 931). *Dominance* requires a reward structure capable of offsetting the subjective costs or values associated with the process of making and executing individual decisions: “The reward structure

dominates any subjective costs (or values) associated with participation in the activities of an experiment” (p. 934). Dominance is intended to ensure that the choices individuals make under the first percept are actually ‘costless’ choices. *Privacy* restricts the information of each subject to his own rewards: “Each subject in an experiment is given information only on his or her own payoff alternatives” (p. 935). Privacy is therefore meant to guarantee the ‘autonomy’ of the decision-makers by forcing them to concern only about their own rewards.^{68,69}

The goal of nonsatiation and saliency is to ensure that individuals act in accordance with the framework of the institution so that individual behaviour and aggregate outcomes can be appropriately tied to the microeconomic institution. They are meant as sufficient conditions for the creation of an experimental microeconomy in the sense that they can make the microeconomy meaningful to the experimental subjects.⁷⁰ They are not sufficient conditions for a controlled microeconomic experiment, however. Dominance and privacy must be verified to ensure that only the factors pertaining to the design of the experiment determine subjects' decisions.

It is easy to see that the four precepts aim at controlling the ‘environment’ of the microeconomic system. They ensure that experimental subjects act in conformity with the circumstances defined by the institution. The ‘institution’ of the microeconomic system is instead controlled by the rules under which experimental subjects are to communicate, exchange or transform commodities.

The four conditions of a valid experiment are deemed sufficient to warrant belief in the experimental results. That is, they are considered sufficient to provide economists with reasons for believing that the observed behaviour is determined by factors pertaining to the experimental design and by these factors only. But they do not allow drawing conclusions about real world microeconomies. To do that, the experimental microeconomy must fulfil a fifth precept: the precept of parallelism.⁷¹ *Parallelism* asserts that “(p)ropositions about the behaviour of individuals and the performance of institutions that have been tested in laboratory microeconomies apply also to nonlaboratory microeconomies where similar *ceteris paribus* conditions hold” (p. 936). These concern the variables that “were constant in the laboratory experiments” (p. 937, n. 20). However, Smith notes that the claim to generality is first and foremost a qualitative claim that pertains to the relation between the environment,

⁶⁴ There are also hybrid experiments that combine aspects of both kinds of experiments (e.g. monopoly market experiments or auction market experiments). Nonetheless, these two categories comprehend distinctive types of experiments each of which is associated with a particular body of theory.

⁶⁵ To put it simply, the surplus is the gain of exchange to buyers and sellers. The buyer's surplus is given by the difference between the reservation price of the good bought and the price paid for it and the seller's surplus is given by the difference between the price of the good sold and its cost of production. The players' strategies in a game constitute a Nash equilibrium when each chooses a strategy that is the best (utility-maximizing) response to the other players' equilibrium strategies. In decision-making under risk, expected utility theory posits that consistent choices are choices that satisfy the conditions of completeness, transitivity, continuity, reduction, independence and monotonicity (see Ch. 9, n. 162).

⁶⁶ These are not intended as necessary conditions for creating and controlling laboratory microeconomies.

⁶⁷ In experimental economics an ‘important’ payoff pays subjects an amount above the hourly wage of the subject pool. A flat non-salient payment is added to the salient amount to ensure that subjects at least receive the minimum payoff.

⁶⁸ Siegel and Fouraker (1960) and Fouraker and Siegel (1963) were the first to demonstrate that intersubjective considerations could be controlled by the condition of incomplete information in which subjects are only knowledgeable of their own payoffs. However, as will become apparent in the remainder of this thesis, the feasibility of implementing this condition depends on the problem the experiment is designed to solve.

⁶⁹ The articulation of the conditions for a controlled experiment was first put forward by Wilde (1981) who qualified Smith's (1976a) prerequisite of nonsatiation with the precepts of saliency, dominance and privacy. However, in Smith (1982) saliency is a demand for a reward structure that links the subjects' payoff with the outcome rather than the actions taken by the subjects.

⁷⁰ Smith notes that the ‘game value’ individuals attach to experimental outcomes may have the same effect as the monetary reward structure. However, the fact that ‘game value’ is another motivational factor requires the unconditional adoption of a monetary reward structure to induce the desired preferences on subjects (1982, p. 935).

⁷¹ The distinction between the establishment of an experimental result and its transferability to non-experimental contexts is captured in the experimental literature as a distinction between the internal and external validity of experimental claims. This topic is scrutinised in depth in the next chapter (section 6.5).

the institution and the performance of the system. For more quantitative results, parallelism would require a tighter correspondence between the laboratory and the nonlaboratory environments (p. 937). Whether or not the percept of parallelism holds can only be determined by additional empirical means.⁷²

THE EXPERIMENTAL MICROECONOMIC SYSTEM

In the remainder of this thesis I show how SEE can be brought to bear on the method, the arguments, the experiments, the results, and the controversies of experimental economics. In order to do that some preliminary terminological clarification is required given that SEE was built upon the study of experimental practice in physics.

In section 2.3, an experimental system was decomposed into three major components. The experimenters' intervention in the material world to setup and operate material apparatuses in the generation of data about the object under scrutiny is the *material procedure* of the experimental system. The *instrumental model* consists in the conceptual understanding of how the apparatus functions and how it yields information about the object under scrutiny. Finally, the *phenomenal model* interprets the experimental data and translates them into meaningful information about the object under scrutiny. Experimental practice was then conceptualised as the forging of a *three-way coherence* between the three components of the experimental system. When coherence obtains experimenters acquire confidence that they have understood their practice and the outcome of practice. This was presented as a very general and abstract framework that accounts for the way in which experimenters produce and establish knowledge via experimental systems purposefully designed and implemented to provide answers to particular problems.

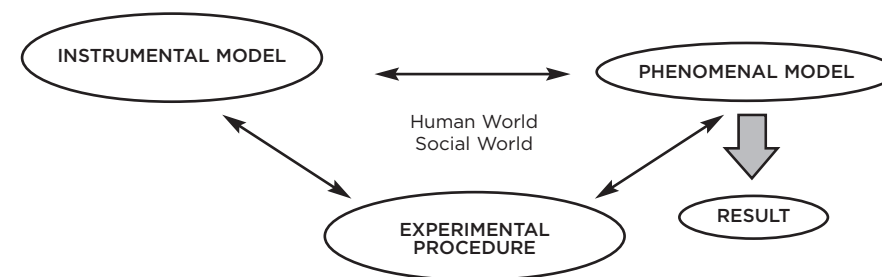
An economic experiment is to be conceived of as an experimental system that is composed by the experimental procedure, the instrumental model, and the phenomenal model. The design and the enforcement of the set of rules that structure the interactions of experimental subjects constitute the *experimental procedure* of an economic experiment. The experiment's instructions, the laboratory room, and the terminal computers, are all part of the 'material apparatus' whereby experimenters generate data about laboratory microeconomies. The *instrumental model* consists in the conceptual understanding of how the experimental microeconomy is to contribute to the acquisition of knowledge about individual behaviour and microeconomic institutions. The interpretation of that behaviour and the institutional performance is in turn the experiment's *phenomenal model*. In an economic experiment, when a *three-way coherence* obtains the experimenters acquire confidence that they have understood the behaviour of the subjects and have made sense of the microeconomic system's performance.

⁷² It should be noted again that this section presents the stylised view of an economic experiment to set the background for the analysis of different kinds of experiments carried out in subsequent chapters. A critical assessment of the methodological views that the stylised conception of experiment conveys or suggests is postponed until Ch. 6.

The 'material world' of a microeconomic experiment consists in its objects of scrutiny, which are the experimental subjects and the experimental microeconomic institutions. Given that we are now in the realm of the human sciences, the expression 'material world' is substituted for the expression 'human world', which is used to refer to both experimental subjects and experimental institutions. By the same token, the expression 'material agency' is substituted for the term 'human agency' to account for the potential resistance of the objects of scrutiny to experimenters' anticipations. In economics, the 'human world' resists whenever observed individual behaviour and their aggregate outcomes frustrate scientists' expectations. This frustration is jointly determined by the characteristics of experimental subjects and the experimental context wherein their actions take place.

Finally, just as 'material agency' is a source of epistemic value in the natural sciences, so is 'human agency' in economic experiments. That is, the epistemic value of economic experiments derives from the capability of the experimental subjects to produce behaviour that frustrates experimenters' expectations.⁷³ Of course, and similarly to the natural sciences, the epistemic value of experimental results also derives from the collective nature of knowledge production, or the economists' interaction with the social world. In the next section the meaning and significance of these notions is exemplified with Smith's first market experiments. For now, the foregoing exposition is synthesised in figure 5.1 below.

FIG. 5.1. The three-way experimental coherence of experimental economics



5.3 THE MARKET EXPERIMENT

Smith's first contact with experimentation goes back to late 1940's when he was a graduate student at Harvard. It all got started when Smith participated as a subject in one of Edward Chamberlin's pedagogical experiments that allegedly demonstrated the predictive failure of the standard neoclassical model of a market under perfect competition. At that time, Smith thought that "the whole exercise was sort of silly".

⁷³ This intuition is already present in extant accounts of economic experiments (e.g. Morgan 2002, 2003, 2005). This topic is addressed in detail in chapter 8.

However, it was impressive enough to catch his attention in one night of insomnia when it occurred to him that “the idea of doing an experiment was right”. What was wrong was Chamberlin’s experimental design. On his view, a more powerful design would have to give the theory of perfect competition a better chance so that its failure would be a more powerful result. These thoughts kept him dwelling on what could possibly be an appropriate design for “the more credible job of rejecting competitive price theory” (1991, p. 370). In 1956, now as an assistant professor at Purdue, he decided to give it a shot. In so doing he initiated the process that would significantly contribute to the foundation of the discipline of experimental economics.

Smith’s *problem-situation* arose from his discomfort with Chamberlin’s experiment and the expectation that it could be improved. A more adequate test to competitive price theory (CPT) would require the design and implementation of a new experiment. Nonetheless, CPT and Chamberlin’s experiment were still the main items of scientific culture Smith resorted to. CPT predicts that a competitive market reaches a state of equilibrium (defined by an equilibrium price and an equilibrium quantity). That is, equilibrium is reached in a market characterised by atomistic buyers and sellers (no one has significant power to affect the terms of trade of a homogeneous good), perfect information (regarding the terms of trade and the good) and perfect mobility of factors (factors can be put to alternative uses with no cost). At equilibrium buyers and sellers can do no better. At the equilibrium price, no buyer wants to buy an extra unit of the good nor does any seller want to sell an extra unit of it or, in other words, supplied quantity equals demanded quantity. Moreover, this is a point where efficiency is maximised because total surplus is at its maximum level.

In Chamberlin’s experiment (1948) students would circulate in the classroom and engage in bilateral bargaining, acting either as buyers or sellers of a fictitious good, until a contract was made or the trading period ended.⁷⁴ As mentioned above, according to Smith a more appropriate setting would have to be more favourable to CPT so that a negative result would be more compelling. To that effect, two important modifications were introduced to Chamberlin’s *experimental procedure*. First, the bilateral bargaining procedure was replaced by a multilateral procedure to better capture the informational structure of competitive markets. The problem, as Smith saw it, was that under Chamberlin’s design there was no opportunity for the dissemination of information, but “perfect” competition is characterised by perfect information about the market conditions, namely about all offers to buy and sell. Smith chose the auction mechanism of the New York Stock Exchange to allow the simultaneous eliciting of quotations for

the whole trading group, rather than to a single trader at a time, and thereby provide traders with information about the market conditions. Thus, all traders could observe the tendered offers and bids and whether they were accepted or not. Second, a multi-trading period set-up substituted the single period design to test the hypothesis that markets tend to approach equilibrium over time (rather than the alternative hypothesis of instantaneous equilibrium). In addition, the multi-trading period set-up would provide students the opportunity to acquire experience, as it was typical of ‘real’ market traders (Smith 1962, n. 5).

Smith’s *experimental procedure* was the following. Subjects were randomly assigned the roles of a seller or a buyer. Each seller was given a card containing a number that represented the minimum price for which he would be willing to sell one unit of a fictitious indivisible commodity (i.e. his reservation price). It was explained that sellers could not sell the commodity at a lower price but could sell it at a higher price and thereby earn the difference between the actual contract price and the minimum reservation price marked on the card (the seller’s surplus). A seller’s reservation price was private information. Similarly, each buyer was assigned a private reservation price representing the highest price at which the unit of the commodity should be acquired. Purchases at a price lower than the reservation price resulted in earnings of the difference between the maximum reservation price and the actual contract price (the buyer’s surplus). Moreover, subjects were told that they should also be willing to make contracts at their reservation prices rather than leaving their wants unsatisfied (or failing to make a sale) (1962, p. 112). By assigning reservation prices for each trader Smith defined the market’s supply and demand curves, i.e. the set of possible supply and demand quantities at each price, and the equilibrium values for quantity and price (obtained at the intersection of the supply and the demand curves, which assume that buyers and sellers are willing to trade at their reservation prices). Because reservation prices were private information, experimental subjects could not possibly compute the equilibrium price.

In each period (out of a total of five periods, each of about five to ten minutes long depending on the number of participants) buyers and sellers would engage in multilateral bargaining for the exchange of a unit of the commodity by crying out loud price offers for buying or selling. Whenever a match was reached a binding contract would be closed for the agreed contract price.⁷⁵ If there were no more offers or the trading period finished, a new trading period would resume and the same conditions would be restored for each subject. The data of the experiment would consist in the reservation prices and the contract prices for each commodity traded during market exchange.

The description of the *experimental procedure* characterises the ‘institution’ and the ‘environment’ of the microeconomic system. It is easy to see that the institution of the microeconomic system determines that the economic agents are to exchange a fictional commodity by orally stating intentions to buy or sell to the whole group of traders, and that this commodity is traded whenever there is a match between a bid and an offer

⁷⁴ Chamberlin’s experiment is a classroom experiment rather than a laboratory experiment (so were Smith’s first experiments). This means that this experiment does not adequately satisfy the precepts of experimental economics. It probably does not satisfy the precept of ‘saliency’ in that subjects’ rewards were hypothetical. Even though subjects did not know each other’s hypothetical payoffs, the condition of privacy might not have been fully satisfied either. This is because subjects were engaged in direct interaction with a classmate, which might have prevented subjects’ ‘autonomy’. In addition, subjects were engaged in a bargaining situation with fellow-students before the eyes of their teacher. In contrast, in laboratory experiments, subjects receive monetary payoffs depending on the subjects’ actions and privacy is guaranteed by preventing subjects from knowing with whom they are interacting. At any rate, Chamberlin himself did not take much notice of these experiments beyond the pedagogical function of demonstrating to students what he thought were the weaknesses of CPT.

⁷⁵ Buyers would state aloud “Bid x \$” and sellers would cry “Offer y \$”, whenever $x = y$ a contract would result.

price. The environment of the microeconomic system is characterized by the number of traders and their reservation prices, which define agents' willingness to buy or sell one unit of the commodity and thereby the overall supply and demand for each price.

The experiment's design is conceived to effect control over subjects' motivations so as to render the experimental situation relevant to them. This is achieved by making sure that subjects understand the experimental task at hand and by being very explicit about what is expected from them. In addition, control is reinforced by the implementation of the precepts of experimental economics. Even though in the first market experiments subjects did not earn monetary rewards, subjects' motivations were believed to be aligned with the required motivations. That is, subjects were deemed to be adequately motivated to maximise fictional payoffs.⁷⁶ In short, saliency was guaranteed by having fictional payoffs dependent on subjects' actions and privacy was guaranteed by making the knowledge of reservation prices private information. Nonsatiation and dominance were guaranteed by the incentive structure of the experimental task that motivated subjects to try to attain the highest fictional payoff.

The new experimental design was conceived of as a better test to CPT. The new experimental market would purportedly provide a better refutation of CPT than Chamberlin's experimental market. The failure of CPT under more favourable conditions would be a far more conclusive result. This was the *instrumental model* that would lead to the refutation of CPT were the experimental market to fail achieving the equilibrium price and quantity – the *phenomenal model* Smith was willing to accept and that would render the experiment meaningful and intelligible.

Even though the experimental set-up was favourable to CPT, Smith noted that there would be room for market failure. For “[t]he mere fact that, by any definition, supply and demand schedules exist in the background of a market does not guarantee that any meaningful relationship exists between those schedules and what is observed in the market they are presumed to represent. All the supply and demand schedules can do is set broad limits on the behavior of the market” (1962, p. 114-5). In other words, Smith is here claiming that even though the experimental procedure controls subjects' motivations (to elicit maximising behaviour), the experimental conditions warrant enough room for *human agency*. The experimental conditions did not force subjects to trade at equilibrium, which was not known to them. The contracts could be closed for a wide range of prices different from the equilibrium price. But, in fact, it did not. Rather than obtaining a categorical rejection of CPT, the experiment produced confirming evidence.

Because the results of the first experiment yielded the opposite phenomenal model – the success of the competitive market in reaching equilibrium – Smith doubted them. The three-way coherence achieved was simply not convincing enough. Full confidence could only be achieved by way of conducting more experiments. And so he did. Smith performed a series of more experiments to check whether or not the

original result could be attributed to a faulty experimental procedure. However, subsequent experiments replicated the original result: the competitive experimental market converges to equilibrium. In 1960 Smith finally decided that he had accumulated sufficient data that demonstrated this result was not an artefact of the experimental procedure.

In 1961 Smith submitted his findings to the *Journal of Political Economy*. Five years of supporting results were necessary to give Smith confidence in his own results, or perhaps more importantly, assurance that the community of economists could also trust them. The reproduction of the original result by various experiments could more easily persuade the community of economists that this result was valid. Smith's willingness to report these experiments was ultimately encouraged by his co-practitioners who persuaded him to publish what he knew was a controversial proposal: the experimental method for economics. This is how Smith recalls it:

Whatever the exact genesis, I got up the courage to write a paper reporting on all the experiments I had done from 1956 to 1960. It wasn't easy. People had been sceptical that there was a trick, some people reason why the experiments worked that had *nothing* to do with economics or theory or that overused, undefined thing that economists call the “real world”. But there were also those who consistently encouraged me – John Hughes and Em Weiler, in particular. (1991, p. 372 emphasis in original)

Besides the support received from his peers, Smith's decision to publish the experimental results might have also been encouraged by the publication, at that time, of other pioneering experimental work. This is the point to which I turn to in the next section.

5.4 THE EARLY NON-MARKET EXPERIMENTS

It is remarkable that the three grand categories of economic experiments emerged with equivalent expression at about the same time (notwithstanding some earlier dispersed breakthroughs). It is no accident that *individual behaviour and decision-making experiments* and *game theory experiments* came out at the same time and grew in tandem. The reason is that they share at least two historical factors that strongly contributed to their emergence: a major theoretical development in economics and psychologists' exportation of the experimental method to economics.

The contribution of John von Neumann's and Oskar Morgenstern's *Theory of Games and Economic Behavior* (1944) to experimental economics is widely acknowledged. It generated an array of novel predictions that were suitable for experimental research, which was already a very familiar method to psychologists at the time. The first incursions in the field of *individual behaviour and decision-making experiments* aimed at testing von Neumann and Morgenstern's assumptions about individual choice. An important contributor was Maurice Allais (1953) whose name became inexorably associated with probably the most well-known ‘anomaly’ to Expected Utility Theory (EUT) – the Allais paradox. The Allais paradox, which

⁷⁶ Smith notes that “the subjects have shown high motivation to do their best even without monetary payoffs” (1962, p. 20). The claim is that the attempt to make the highest fictitious profit was a sufficient incentive for subjects to act in conformity with the experimental context. This is supported by the results obtained that revealed that subjects did attempt to maximise the fictional profits.

is associated with ‘the common consequence effect’ and ‘the common ratio effect’, caught the interest of economists to the experimental study of the violations of EUT. Ever since individual choice experiments have been prolific in collecting systematic violations to individual decision-making theory and a most inspiring field for the exploration of new theories.⁷⁷

Initially, the *game theory experiments* aimed at testing hypotheses about the feasibility of stable equilibria in various strategic problems. The earlier references include Melvin Dresher and Merrill Flood’s test of what subsequently came to be known as the prisoner’s dilemma (Flood, 1958), the n-person games by Gerhard Kalisch, J. Milnor, John Nash, and E. D. Nering (1954) and the coordination games by Thomas Schelling (1957). Game theory experiments have since uncovered systematic violations of theory and have been also an important source of theoretical development. Theories of other-regarding behaviour that incorporate non-economic factors such as fairness considerations and the evolutionary models of learning and adaptive behaviour are among some of such examples.⁷⁸

Even though these experiments were carried out more or less independently, these early works had a subsequent impact on one another. Perhaps the most significant influence was methodological in that this early work clearly set the standards for the practice of experimental economics. The work of Fouraker and Siegel was, in this regard, of particular importance in setting three requirements that still characterise current practice: the reporting of experimental instructions to promote the replicability of experimental results by others; the enforcement of anonymity to ensure effective control over subjects’ motivations; and the use of monetary incentives to induce economic behaviour. This is how Smith acknowledges his debt to Sidney Siegel:

Sid Siegel’s concern for the content of instructions, his insistence that instructions should be part of one’s report, his practice of always motivating subjects with meaningful cash rewards, and much more are all evident in this work [Fouraker and Siegel’s 1963], in his other publications, and in the work of experimental economists today. (Smith 1992, p. 247 emphasis omitted)⁷⁹

Even though this does not emerge in the historical accounts of experimental economics, the independent appearance of experimental economics at about the same time seems to have been inspired by the broader context of research carried out under the banner of RAND Corporation.⁸⁰ In the late 1940’s, right after World War II, RAND Corporation re-oriented its research to peacetime selecting the field of

‘cybernetics’ as the privileged area of investigation. The goal was to further the application of mathematics to a wide range of human affairs beyond warfare applications. The main areas of research were computer science, artificial intelligence, information theory, mathematical economics, mathematical learning theory, game theory, automata theory and operations research, among others. This context was important to the development of game theory, and perhaps more fundamentally, it provided the opportunity for economists to contact and become aware of the potential contribution of the experimental method to economics. This was indeed facilitated by the highly interdisciplinary and interactive context that RAND Corporation provided. It brought together physicists, mathematicians, statisticians, economists, psychologists among other natural and social scientists. In this connection it is useful to cite the most explicit statement by one prominent witness, Herbert Simon:

I do not think that the impetus for experimentation within a game-theoretical framework initially came from economists, but rather from psychologists (particularly those who had begun to build mathematical learning theory), statisticians, and interdisciplinary types close to cybernetics and management science... Siegel, who I always supposed was the key member of the Siegel/Fouraker team, was a psychologist. Experiments, of course, came naturally to psychologists – that is the way they learned to do science. (in Smith 1992, p. 253-4)

It is remarkable that all of the first contributors to the field were researchers at RAND Corporation or had contact with the ongoing research through various means, such as in seminars and conferences.⁸¹ The 1952 Santa Monica Seminar “The Design of Experiments in Decision Processes” is a prominent example of just that. Indeed, the conference proceedings *Decision Processes* (Thrall *et al.* 1954) contain most of the seminal contributions to experimental economics. This influence continued in subsequent developments, as will become apparent below. For now it suffices to mention the influence of RAND on the Cowles Commission and, in particular, on the work of the welfare economists Tjalling Koopmans, Kenneth Arrow and Leonid Hurwicz, who contributed to the theoretical development of another fertile area for the use of the experimental method: social choice theory.^{82,83}

⁸¹ Indeed the list of the researchers who collaborated with RAND during this period is quite revealing: Melvin Dresher, Merrill M. Flood, John von Neumann, Thomas Schelling, Herbert Simon, and, of course, Vernon Smith.

⁸² The Cowles Commission for Research in Economics was founded in 1932 by Alfred Cowles at Colorado Springs, USA. It turned into the Cowles Foundation for Research in Economics, in 1955, when it became an institute of the Department of Economics of Yale University. Back then, the major research goal was to foster the development of logical, mathematical and statistical methods of analysis for their application in economics and other social sciences. In the post-war period, in particular, the broader aim was to improve state intervention and guidance in economic and social issues (see Mirowski 2002, Ch. 5 and <http://www.cowles.econ.yale.edu>).

⁸³ In an extensive book dedicated to the study of the influence of post-war military institutions on economics, Philip Mirowski (2002, Ch. 8) identifies the mark of ‘cyborg science’ practiced at RAND Corporation on experimental economics. Mirowski identifies as elements of cyborg science: the ‘mechanical rationality’ imposed on experimental subjects, the design of ‘machine-like’ market institutions, the computerisation of the experiments, the substitution of subjects by simulation programs, and the treatment of social contexts as games between isolated asocial subjects. In the remainder of the chapter no further references will be made to cyborg science. However, Mirowski’s depiction of experimental economics as a token of cyborg science will become apparent in the subsequent sections of this and in subsequent chapters.

⁷⁷ For a survey on this experimental work see Camerer (1995).

⁷⁸ See Roth (1995a) for an historical account of experimental economics.

⁷⁹ In his 1962 article Smith quotes Siegel and Fouraker (1960). See Smith (1992) for other first-person accounts about the mutual influence pioneers exercised upon one another.

⁸⁰ RAND Corporation is a non-profit organisation that resulted from the process of separation and independence of Project RAND from Douglas Aircraft Company of Santa Monica, California, USA, in 1948. The name results from the contraction of the term research and development and it became an independent organisation to address the need of connecting, in the cold war period, military planning with research and development issues. The early concern was with defence-related issues and the guiding principle was that domestic defence required more than technology research and development to encompass the direct participation of government in broader social and economic issues (see <http://www.rand.org>).

This brief historical incursion into the birth of laboratory experimentation in economics is very illustrative in showing the importance of interdisciplinary work to the emergence and development of new areas of research. Interdisciplinary work encourages the adoption of foreign methods, concepts, theories, and instruments in new domains and thereby contributes to the discovery and development of unforeseen fields of research. In addition, it adds epistemic value to those discoveries insofar as the imported resources had already been tried out and established in previous practice. In the remainder of this thesis further examples and illustrations are provided of other interdisciplinary exchanges that contributed to the development of the field of experimental economics.

5.5 ARGUING FOR THE EXPERIMENTAL METHOD

In April 1962, Smith's pioneering results were published in *The Journal of Political Economy*. The amount of time it took until Smith finally decided to publish the results of market experiments cannot be explained by the controversy surrounding the experimental results. The results obtained provided supporting evidence for a well-established theory. The problem was that the experimental data was not considered very informative. First of all, the idea of an economic experiment was simply foreign to most economists. But even if economists had been willing to consider the experimental method, Smith's results would have been deemed uninteresting. The fact that Smith's experiment created favourable conditions to competitive markets rendered the experimental data weak confirming evidence for CPT. Consequently, Smith's market experiments are more adequately interpreted as a justification for the use of the experimental method, rather than a defence of the rather shallow results.

Smith was very cautious in deriving implications from his first set of experiments. The market experiments were presented as tentative testing devices which needed further experimental work. The experimental market was described as a 'simulation' of a real market that attempted to capture essential aspects of markets favourable to CPT. The data generated constituted evidence for the success of CPT in predicting equilibrium in experimental markets. This was demonstrated by the observation that "as the exchange process is repeated through successive trading periods with the same conditions of supply and demand prevailing initially in each period, the variation in exchange prices tends to decline, and to cluster more closely around the equilibrium" (1962, p. 116). No further implications were derived. The theory simply predicts well in these markets.

Smith was, however, more concerned with explaining the concept of an economic experiment. To that end the experimental design and procedure were carefully described (cf. section 5.3 above) with the help of the conventional supply and demand charts depicting the crucial aspects of the experiments. The experimental data (i.e. the experimental contract prices for each trading period) were presented with the use of diagrams. This is rhetorically ingenious. After all, Smith's experiments are all about markets, supply and demand curves, and equilibria – the key concepts of economics!

It is clear that Smith felt the need to build a strong coherent system in support of both experimental results and the experimental method. Smith reported ten experiments, not only to demonstrate that he had collected sufficient evidence that supported his main result, but also to show the legitimacy of the experimental method. He did so by supplying further evidence that the result is robust to altering conditions of the experimental design. Indeed, the results were reproduced under different supply and demand conditions (changes in the position and shape of supply and demand schedules), different levels of subject experience, different market mechanisms (auction and retail market), volume of trade, among other modifications (subject pools, location, etc.). Hence, Smith's support for his data and for the experimental method is grounded on various *three-way coherences* obtained with different experimental set-ups each of which yielded the same result: convergence of the experimental market to equilibrium. To put it in another way, this result was epistemically strong because it was supported by various experimental systems. This strength, in turn, argued for the adequacy of the experimental method.

But these were not the only coherences Smith produced and invoked to justify his results and the experimental method. The results were also sustained by a coherent resolution between theory and evidence. Not only is CPT supported by evidence, but the effects of altering conditions were also correctly predicted by theory (e.g. decreases in demand reduce market prices). Moreover, the statistical tests carried out after the data had been collected rejected the null hypothesis of instability, thus providing further assurance that the data constituted significant evidence for the equilibrating tendencies of the experimental markets.

The series of experiments produced data that could be used to test additional theoretical hypotheses. For instance, it generated data that was used to test the effect of the relative magnitude of buyers' and sellers' surpluses on deviations to competitive prices. Smith also found strong support for the conjecture that the speed of adjustment is higher, the higher the surplus of the traders because it decreases their bargaining resistance.⁸⁴ Moreover, it was found that the relative magnitude of the buyers' and sellers' surpluses affects how the actual market equilibrium is approached. The market equilibrium tends to be approached from below the equilibrium price when the consumer surplus is superior to the producer surplus, because it increases the producers' bargaining resistance, and vice-versa. However, this resistance tends to weaken as traders realise that the envisaged surplus is unattainable.

Coherence was not absolute, however. One out of ten experiments did not support the equilibrating tendency of the market. But this did not pose a special problem because this was a result that Smith could understand and explain. The result occurred under very special conditions that intended "to simulate an ordinary retail market", i.e. a market characterised by an asymmetry between buyers and sellers where only sellers were allowed to make price quotations. Buyers could only accept or reject the offers of the sellers. Under such circumstances the disequilibrating tendency operated in favour of the buyers. Smith explained this by the possibility of concealing their eagerness to

⁸⁴ This was dubbed the 'excess-rent' hypothesis. The other hypothesis tested was the Walrasian hypothesis that posits that the growth rate of exchange prices varies positively with the excess demand at each price.

buy. By explaining the disequilibrium, Smith undermined the significance of the conflicting result. In fact, it contributed to reinforce the general conclusion since this experimental market deviated from the theoretical competitive market. Thus, the apparent destabilisation also supported the consensual experimental result. These results taken together allowed Smith to draw the following conclusion:

[I]n static equilibrium a competitive market minimizes the total virtual rent received by buyers and sellers. By “virtual rent” I mean the rent that would be enjoyed if all buyers and sellers could be satisfied at any given disequilibrium price... [therefore] it is completely plausible to require, in the interest of strict market efficiency, that no trader be imputed more rent than absolutely necessary to perform the exchange mechanics... Hence, at price P [higher than equilibrium price]... virtual rent exceeds equilibrium rent, and if this price persists, some seller get more rent than they “should”. (1962, p. 132-3)

In other words, Smith is claiming that the market experiment has shown yet one more virtue of the competitive market. It prevents “the “wasteful” payment of excessive economic rent” (op. cit., p. 134).⁸⁵

The analysis of Smith’s experimental practice illustrates well the difficulties inherent to scientific practice in new fields of research. In unexplored fields a more complete set of arguments must be provided in support of novel results and of the means by which the results were generated. Similarly to experimental physics (cf. section 2.5), these arguments rely on the same coherence procedures that help constructing a strong net of relations of coherence. The experimental results of economics were supported by: the use of different experimental systems, the confirmation of theoretical predictions, the explanation of anomalous results, the use of data processing techniques, tentative interpretations of the results and their test, and last but not the least, the consistency with a major item of economics culture – competitive price theory.

5.6 THE MARKET EXPERIMENT: THE EPISTEMIC APPRAISAL

But what can we make of Smith’s practice? As Smith presented, the market experiment was designed with the expectation that it would generate negative evidence for CPT. The idea behind it was that negative evidence obtained under favourable conditions to the hypothesis under test is far more compelling than disconfirming evidence obtained under adverse conditions. But the market experiment produced positive evidence to CPT which raises the question – *What is the import of confirming evidence produced under favourable conditions to the theory under test?* The answer to this question is provided via the application of the four epistemic tests of experiment: the materiality test, the stringency test, the technological test and the social robustness test.

SCORE OF THE MATERIALITY TEST: LOW

As mentioned earlier, the concept ‘materiality’ has a special meaning in economic experiments. The ‘material quality’ of the experimental system and the participation of ‘material agency’ in the process of knowledge production, which the materiality test assesses (see section 4.2), must also be translated into the language of experimental economics. On this view, then, the ‘material quality’ of economic experiments refers to the ‘socioeconomic quality’ and to the ‘human quality’ of the three component of the experimental systems: the object of analysis, the experimental practices and the outcome of practice. Similarly, the ‘material agency’ of economic experiments refers to the ‘socioeconomic agency’ and to the ‘human agency’, that is, the contribution of the microeconomic institution and that of subjects’ actions to the generation of the experimental outcomes.

The score of the materiality test of the market experiment is low because the socioeconomic and the human ‘quality’ and ‘agency’ of the experiment are low. On the one hand, the experimental market significantly deviates from ‘real’ markets and, on the other hand, it strongly determines the behaviour of the experimental subjects. To put it in another way, the low level of materiality of the experimental market is explained by the high ‘theoretical’ content of the microeconomic institution and the high degree of control exercised over subjects’ behaviour.

The experimental market is more adequately perceived as a model of CPT rather than a model of a real market. The reason is that the experimental market was built upon CPT’s conception of a market. Information about real markets was only resorted to when the theory revealed to be an insufficient guide for the design and material implementation of the experiment. Information about ‘real’ markets was only used to help implementing the theoretical competitive market in the laboratory. As Smith put it, the experimental market is “a good general model of received short-run supply and demand theory” that simulates the auction-trading process of the stock, bond, and commodity markets, because they provide “the best chance of fulfilling the conditions of an operational theory of supply and demand” (1962, p. 111). So the fact that the experimental market did not conform to what was known about many kinds of real markets was not a problem. For instance, the greater constancy of the conditions of the experimental markets (in relation to real markets) was not problematic because stability was needed to study the equilibrium tendencies of competitive markets. For equilibrium is “a condition toward which the market would move *if* the forces of supply and demand were to remain stationary for a sufficiently long time” (op. cit., p. 115, emphasis in original), and it was this concept of equilibrium that the experiment was designed to study. Thus, the experimental market is better perceived as a model of a theoretical market rather than a model of a real market. Moreover, the experimental market is made of students performing the role of buyers or sellers (rather than real traders) who engage in the exchange of a fictitious good (rather than a real commodity). In sum, the socioeconomic and the human ‘quality’ of the experimental market is reduced insofar as it is more adequately perceived as an instantiation of a theoretical market while it differs in significant ways from ‘real’ markets. Nonetheless, the market

⁸⁵ In the appendix Smith provides further analysis of the data, including: inter-trading period, the convergence path of bid, offer and contract prices, and a comparison between the short-run and the long-run equilibrium.

experiment still managed to implement in the laboratory a special market in which experimental subjects engage in the exchange of a fictitious good.

The human 'agency' of the market experiment is also reduced. On closer inspection, one observes that the market's trend to equilibrium is very strongly determined by the design of the experiment. The design of the experiment severely constrained subjects' decisions by being very effective in inducing on the subjects the behaviour assumed by CPT. Subjects were explicitly told to behave as theoretical sellers (or buyers) by receiving explicit instructions how to maximise their earnings (by selling at high prices or buying low) and by being prevented from trading with a loss (selling below reservation prices or buying above them). The enforcement of the precepts of experimental economics, in addition, gave them no reason to do otherwise. Thus, it is not surprising that subjects' behaviour conformed to the theory's predictions.

In addition, and most crucially, the structure of the experiment was very effective in generating a trend towards equilibrium for reasons other than the characteristics of the microeconomic institution. Equilibrium is explained by the pattern of values induced on the experimental commodities and the reposition of the market conditions at the beginning of each period. In this experiment, the 'commodities' with higher redemption values are most easily purchased because they are the most profitable and therefore they can be purchased at higher prices. Similarly, the 'commodities' with lower production costs are most easily sold with a profit at lower prices. These two factors taken together imply that after the initial contracts, the tradable commodities remaining in the market are those with lower redemption values (which can only be bought at lower prices) and with higher production costs (which can only be sold at higher prices). Consequently, it is not surprising that as the trading period elapses the remaining goods are transacted at prices that approximate the equilibrium price (which is in between the remaining highest redemption values and lowest production costs). That is, as the trading period progresses the subsequent transactions tend to be made on terms that slowly approximate the equilibrium price. The goods with the lowest redemption values and those with higher production costs are gradually excluded from the market. Therefore, not only does the price tend to move towards the equilibrium price but the quantity also tends to the equilibrium quantity.

Moreover, the transaction prices of a period tend to affect the terms of trade of the subsequent period. That is, the last transaction prices of a given period tend to be close to the first transaction prices of the subsequent period. This is so because buyers with high redemption values learn that there are sellers willing to sell at much lower prices and so they try to transact at those prices. The low cost producers, in turn, try to sell at the higher equilibrium price. The price trend is thereby extended from one trading period to another until competitive price is eventually reached and efficiency is maximised.⁸⁶ This trend is accentuated by the stationary repetition of the market conditions at the beginning of each period.

⁸⁶ The strong correlation between the prices of the last transactions of a period and those of the first transactions of the subsequent period was identified in the transaction paths of many repeated double auction experiments (Roth 1995a, n. 88).

The claim that the experimental procedure strongly determines market outcomes has received additional support by the famous simulation by Dhananjay Gode and Shyam Sunder (1993) which demonstrated that even a market with zero-intelligent traders, who were only restricted to making nonnegative profits, succeeds in achieving equilibrium. A success that is explained by the parameters of the experiment that promotes efficient trading (cf. Holt *et al.* 1986 and Banks *et al.* 1989).⁸⁷

In conclusion, the low level of socioeconomic and human 'quality' renders the experimental market a very special institution that constrains 'human agency' to the point that it makes virtually impossible for subjects to act otherwise. By constraining the 'resistance' of 'human agency', the experimental procedure had a great effect in determining the experimental outcome. Because the characteristics of the experimental market depart in significant ways from real markets, the results it produced cannot constitute sound supporting evidence to CPT.

SCORE OF THE STRINGENCY TEST: HIGH

The stringency test of the market experiment is high since this experiment has the potential for producing conclusive results. The experiment was deliberately conceived to assess whether or not market equilibrium is reached. And equilibrium would admit only a few results within the range of admissible outcomes. To be sure, the market experiment is an experimental system that would easily discriminate between the two alternative phenomenal models in confrontation: equilibrium vs. non-equilibrium. Even though the equilibrium was a surprising result from the viewpoint of Smith's expectations, as he expressed them, there was no doubt that this was the resulting phenomenal model.⁸⁸ The stringency of the market experiment is further reinforced by the robustness of the results to modifications introduced to the experimental procedures. This provides further evidence for the extremely high level of control the experimental procedure exercised over subjects' behaviour. In sum, the market experiment is a stringent experimental system that produces transparent and robust results.

SCORE OF THE SOCIAL ROBUSTNESS TEST: HIGH

Economic experiments possess a potentially high level of social robustness because they are easily exposed to public scrutiny. In contrast to physics, where the public appraisal of experiments cannot extend much beyond the community of specialists, in economics both the practice and the outcome of practice are accessible to other experimenters, non-experimenters and non-economists alike. The reason for that derives not only from the simplicity of economic experiments,⁸⁹ but also from the

⁸⁷ For a complete survey on market experiments see Charles Holt (1995).

⁸⁸ Notwithstanding Smith's widely expressed prior beliefs, one does wonder whether these were his real expectations. Allegedly surprising results are rhetorically more effective because they allude to the epistemic superiority of falsified expectations. Conversely, the attainment of an expected outcome is more easily prone to charges of confirmation biases.

⁸⁹ As will be explained at length in the next chapter, simplicity is a value of experimental economics insofar as subjects must understand the experimental tasks in the intended way so that the experimenters can interpret the results accordingly.

nature of the experimental procedure. The experimental procedure of economic experiments consists to a great extent in the set of instructions given to the experimental subjects that are published along with the report of the experimental results or made available upon request.⁹⁰ These two factors taken together render both the practice and the outcome of practice suitable to the scrutiny of a wide and heterogeneous community of researchers. (In fact, this feature is what makes the present work possible.)

In what concerns the market experiment, in particular, the hostile environment Smith encountered further enhanced the level of social robustness. The anticipated and the actual hostility faced towards experimentation forced Smith to adduce a varied set of arguments in support of the market experiment and of the experimental method. In the remainder of this chapter, further strategies are presented that show how Smith continued his endeavour of justifying the use of the experimental method in economics.⁹¹

SCORE OF THE TECHNOLOGICAL TEST: HIGH

The appraisal of the technological success of market experiments is postponed to the next section where I review subsequent developments. For now it suffices to note that the market experiment turned out to be an adequate tool for subsequent practice and that this adequacy resulted from its low level of materiality and high level of stringency. In other words, because the market experiment was a well-understood microeconomic system that was effective in exercising control over human behaviour, it became a useful device to the same effect in the resolution of other problem-situations.

5.7 FROM TEST-DEVICE TO 'TESTBED'

Smith's understanding of the market experiment evolved with his experimental practice. The market experiment turned out to be an inadequate tool to test CPT. However, its uses grew in the course of practice until it became a device for the design of market institutions, experimental and otherwise.

THE 'TECHNOLOGICAL' APPLICATION OF MARKET EXPERIMENTS

Initially, Smith was concerned with providing supporting evidence for the tendency of experimental markets to converge towards competitive equilibrium. This trend was robust to modifications introduced to the shape of demand and supply curves and to the distribution of surplus between buyers and sellers (1964, 1965). With the accumulation of experimental data, further evidence was provided that showed that, contrary to theory, neither complete information nor large numbers of traders is a necessary condition for convergence to competitive equilibrium. As Smith put it:

⁹⁰ Smith's instructions of the first market experiment are available in Smith (1964).

⁹¹ The social robustness of experimental economics is developed in more detail in Chs. 6 and 9.

There are no experimental results more important or more significant than that the information specifications of traditional competitive price theory are grossly overstated. The experimental facts are that no double auction trader needs to know *anything* about the valuation conditions of other traders, or have *any* understanding or knowledge of market supply and demand conditions, or have any trade experience (although experience may speed convergence) or satisfy the quaint and irrelevant requirement of being a price 'taker' (every trader is a price *maker* in the double auction. (1976b, p. 57, emphasis in original)

Having established robust results for the original market institution (which became known as the double auction), Smith set out to investigate other modes of market organisation such as the discriminative sealed-bid and the competitive sealed-bid auction markets (1967), seasonal markets (with Miller and Plott 1977), and their relative performances (Smith 1976b, 1980b; Plott and Smith 1978; Smith *et al.* 1982; Smith and Williams 1983). Having found that the double auction is the most efficient institution, from a deficient test device of CPT it became, with more propriety, a proxy of a competitive market that could be used as a benchmark, i.e. a reference against which other market institutions could be designed and compared with. In Smith own words:

The motivation for selecting the double auction for extensive experimental study (Smith, 1962) was based on the conjecture that this institution was the one under which classical supply and demand theory had the best chance of being validated. However, I did not seriously expect competitive price theory to be supported by these initial probes... But if competitive theory had validity under the double auction, this would provide the "control case" or reference institution against which other forms of market organization could be compared. (1976b, p. 47)

In the course of practice, experimentation with markets became a field of experimental economics, which may now be decomposed into three main sub-areas of research.⁹² To put it briefly, the *field of industrial organisation* studies the factors that affect competition and that give rise to antitrust phenomena such as predatory pricing or collusion (e.g. Smith 1981a, 1981b; Coursey, Isaac and Smith 1984; Isaac and Smith 1985).⁹³ The *field of auctions* studies the comparative efficiency of auction mechanisms for the transaction of special goods (e.g. Smith, 1967, 1976b, 1980b).⁹⁴ Finally, the *field of capital and asset markets* addresses the specificities of the financial markets and related phenomena, as for example the formation of market bubbles (e.g. Smith *et al.*, 1988).⁹⁵

The experimental market became a resource of practice of wider application. It was first perceived as a preliminary tool that could provide "a rigorous testing of our ability to model elementary behavior before confronting such models with field data" (Smith 1976b, p. 62). The preliminary nature of these tests presented the

⁹² This classification follows Kagel and Roth's (1995).

⁹³ See Holt (1995) for a review.

⁹⁴ See Kagel (1995) for a review.

⁹⁵ See Sunder (1995) for a review.

important advantage of guaranteeing theories' independence from field-data during theoretical construction and prior to their submission to field-data testing. The idea was that, in contrast with field-data, "it is never tautological to modify the model in ways suggested by the results of the last experiment" (Smith 1976a, p. 274). There is always the possibility of conceiving another experiment to test the modified model before the model is finally confronted with field data. But soon the experiment became the privileged empirical device that provided the best testing conditions.

As the results from experiments accumulated and found increasing support, they added up to the field's 'stylised facts', i.e. empirical propositions well confirmed in the laboratory. When these facts were brought in closer articulation with theory, they became empirical evidence for the field's theoretical findings. For instance, the first market experiment became the double-auction experiment that generated the stylised fact that "allocations and prices converge to levels near the competitive equilibrium... prediction. This convergence is rapid, occurring in three to four trading periods or less when subjects are experienced with the institution" (1982, p. 943-52). To summarise, in the course of experimental practice, early experiments became well-understood devices that could be identified by their designs (e.g. *the Double Auction*, *the English Auction*, etc.). With each experiment was associated a set of stylised and significant facts (e.g. the First-Price Sealed Bid Auction generates higher prices than the Dutch Auction).⁹⁶ In short, they became the fields' items of scientific culture.

As understanding of the new method grew, experimenters began to recognise the use of market experiments for applications beyond academic purposes. They came to be perceived as adequate policy tools that could provide "empirical justification and preliminary experience for the design of field experiments" and a means of "exploring the policy implications of new institutions, or alterations in existing institutional rules" (1976b, p. 62).

The adoption of the market experiment as a policy tool was stimulated by the context of increasing market deregulation in the 1980's and the resulting need of conceiving new market arrangements for a number of sectors in restructuring. At first, experimental studies were commissioned for policy guidance. They provided suggestions about the best institutional make-up for special markets. For instance, the American Civil Aeronautics Board commissioned an experimental study to David Grether, Mark Isaac and Charles Plott (1989) to evaluate the overall impact of the intensifying competition and deregulation on the aviation industry. The goal was to identify the institution that provided the efficient allocation of the airport capacity while benefiting from the overall contextual tendencies of growth, competition and deregulation. Another study was commissioned by the Federal Energy Regulatory Commission and the Energy Information Administration to Kevin McCabe, Stephen Rassenti and Vernon Smith (1989) to investigate the

properties of an auction market for the sale and transportation of natural gas. However, these commissioned reports provided no more than recommendations that were followed at the will of the patrons.

Only recently did experiments come to be used for actual 'institutional engineering'. They became 'testbed' experiments, i.e. "a working prototype of a process that is going to be employed in a complex environment" (Plott 1997, p. 605). This was the case of the bid-auction commissioned by the Federal Communication Commission to allocate the licenses for wireless Personal Communication Systems, in 1993 (I come back to this experiment below).

Finally, experiments came also to be used to shed light on real world phenomena. To give an example, in 1971 the Atlantic Richfield Company claimed that the exploitation of oil leases in the Gulf of Mexico yielded unexpected low rates of return and this was to be attributed to the 'winner's curse': the winning bidder is the trader who mostly overestimates the value of the oil tract and therefore incurs in high costs (cf. Capen *et al.* 1971). This occurs because these leases are transacted in an auction (known as a common value auction) where the value of the item transacted is unknown at the time of the sale. Under these circumstances, bidders must trade on the basis of their private estimates of that value. Consequently, the winner of the auction is the trader who possesses the highest estimate which is an excessively optimistic one.

The claim that the winner's curse was occurring in the exploitation of the oil leases in the Gulf of Mexico could only be settled by experimental means. On the one hand, the explanation put forward by the oil company was in stark disagreement with economic theory. While the former claimed that experienced bidders persistently make mistakes notwithstanding the losses they make, the latter predicted that the winner bidder is able to make normal profits because he takes into account that he has the highest private estimate.⁹⁷ On the other hand, field data from common value auctions were inconclusive in that they did not provide relevant information to settle the dispute, namely the private valuations of the transacted items. In addition, the company's interest in convincing other companies to act as a cartel (by reducing their bids) damaged its credibility. In the face of such controversy, John Kagel and Dan Levin (1986) designed an experiment to test the hypothesis put forward by the Atlantic Richfield Company which indeed reproduced the phenomenon in the laboratory.⁹⁸

It is now easy to show that the field of market experiments possesses a high degree of *technological success*. After two decades of practice, the market experiment turned out to be a useful tool for solving various scientific problems. These ranged from empirical testing and the explanation of real world phenomena to 'institutional engineering'. This technological success owes much to the possibility of exercising a high level of control over human behaviour. To put it in another way, the high level of technological success of these experiments hinges on their low level of materiality and high level of stringency. To show this, a closer inspection of a successful story is in order.

⁹⁶ In a First-Price Sealed Bid Auction the commodity is awarded to the highest bid after all bids have been collected; in an English Auction the commodity is awarded to the last bidder remaining in the auction after all other bidders have drop out as the price of the commodity increases; in the Dutch Auction the commodity is awarded to the bidder who accepts to buy the good at the auction price as the price decreases.

⁹⁷ For a review of the winner's curse phenomenon see Thaler (1988b).

⁹⁸ For a detailed analysis of the Kagel's and Levin's experiment see Guala (2001).

A TESTBED MARKET EXPERIMENT

The auction for the trade of licenses for wireless Personal Communication Systems is the most illustrative case of the technological success of market experiments. This design, commissioned by the Federal Communication Commission (FCC hereafter), was purposively studied and tested in the laboratory to be implemented in the field. This was a ‘testbed’ experiment that involved the work of prominent game theorists and experimenters, such as Robert Wilson, Paul Milgrom, Charles Plott and Preston McAfee. The goal was to design an auction that guaranteed efficient allocation of airwave spectrum rights by guaranteeing the sale of the licenses to those who most valued and therefore could make best use of them, and that minimised the cost of the transaction (in terms of time, transparency, etc.).

The auction designed and later implemented in the field was a continuous ascending auction in which traders tendered increasing bids in successive rounds until only one bidder remained. The experimenters were particularly concerned with the presence of opportunistic behaviour, such as the acquisition of licenses for reselling. This opportunistic behaviour was avoided by a set of rules, namely the imposition of an eligibility condition (which required the constitution of a deposit at the beginning of the auction), the enforcement of bidding in the first round, inflicting costs on withdrawals and controlling for the time-length of the experiment so as to avoid the strategic slowing down of the auction.⁹⁹

Before its actual implementation in the field, these rules were tested in the laboratory under conditions that closely resembled the market to be implemented in terms of number of traders, number of items to be traded, complementarities between items, their values, number of rounds, etc. Because the data collected from the lab and from the real auction were similar in many relevant aspects (e.g. bidding patterns, price trajectory, license aggregations, and so forth) the experimenters were confident that an efficient allocation had also been achieved in FCC auction. As Guala put it:

The strength of Plott’s argument [that justifies the efficiency of the FCC auction] lies in the work he and other consultants did to ensure that the same processes took place in reality as those they had observed in their laboratory. The same causes are supposed to operate because experimenters built the two systems so as to be structurally similar to one another. The transportation of the mechanism outside the laboratory was as smooth, gradual and carefully as possible. (2001, p. 473)

It follows from this that the success of ‘institutional engineering’ depends on enforcing rational behaviour in a particular context of interaction and this is attained by finding the ‘right’ institutional framework:

[I]t is by designing and implementing an adequate mechanism that the engineer ensures that rational choice models can work. Since it is partly in virtue of the structure of the situation that economic agents behave rationally, a great part of economic engineering is devoted to ensuring that the structure is ‘right’. (ibid)

To put it in another way, the success of ‘institutional engineering’ relies on the feasibility of designing in the laboratory microeconomic institutions that are capable of controlling human behaviour in the intended way. Moreover, this control must be robust to variability in human motivations and cognitive capabilities so that the same observations obtain with different subject pools. Finally, the technological success of the market experiments also depends on the extent to which the control exercised in the laboratory can be equally exercised in the field (this topic will be further explored in Ch. 6).

In conclusion, the epistemic value of the double-auction experiment increased in the course of practice, as it became a better-understood item of the culture of experimental economics. From a faulty test device to CPT it became a proto-type of an efficient market mechanism which could be used for the development of other market mechanisms. The field of market experiments, in the whole, demonstrated the relevance of institutional arrangements to market outcomes. This is undoubtedly a major contribution of this field of experimental research in economics.¹⁰⁰

5.8 THE NON-EXPERIMENTAL ARGUMENTS FOR THE EXPERIMENTAL METHOD

As mentioned above, it took about six years before Smith ‘got up the courage’ to publicise his results. Smith’s hesitation is explained by the deep-rooted image of economics as a non-experimental science, an image that lasted until at least as late as 1985 (see Appendix 1). It is not surprising, then, that when Smith presented his first results in informal discussions with his peers or in more formal contexts such as at faculty seminars and workshops, he was confronted with scepticism that challenged not only his results but also the whole experimental enterprise. The scepticism encountered was so strong that he was compelled to combine his experimental practice with what could be considered more ‘credible’ work.

In the meantime I had to make a living and during these formative years I was writing and publishing on other topics – capital and investment theory, corporate finance theory and the economics of uncertainty, and natural resource economics. Compared with the experimental work, this other research was much easier to do and easier to publish... The profession was hungry for new theory in these areas, indeed any area that employed recognizable methods;

⁹⁹ Like any other experiment, this ‘testbed’ experiment relied on available theoretical, practical and experimental knowledge. The resources of this experiment included game-theoretical accounts of auction mechanisms (Vickrey 1961; Wilson 1977; Milgrom and Weber 1982), the accumulated experimental knowledge about auction markets, and knowledge about real mechanisms in related markets. See Guala (2001) for their relative contribution.

¹⁰⁰ Even though Smith’s major contribution was in the field of the market experiments, I do not intend to give the impression that this was Smith’s only experimental practice. Smith made also incursions in other topics of research such as public goods (1979a, 1979b, 1980a), bargaining theory (with Hoffman, McCabe and Shachat 1994, and with Hoffman and McCabe 1995, 1996a, 1996b), and individual decision-making (with Knez 1987).

it was not hungry for evidence, and certainly not laboratory evidence, because theory provided all the requisite understanding. (Smith 1991, p. 3-4)

Of course, Smith not only encountered resistance. He was also able to establish alliances with other scientists who shared his interest in the new method. One important ally was Charles Plott, a young colleague that arrived at Purdue in the late 1960s. Plott would later give a substantial contribution to the field, and in particular, to the expansion of the domain of application of the experimental method to social choice theory, public economics and much of political science (Plott 1979; Fiorina and Plott 1978; Plott and Levine 1978). Other important collaborators were Arlington Williams, James Cox, Kevin McCabe, Stephen Rassenti, among many others, with whom Smith established new partnerships when he moved to the University of Arizona, in 1975. In 2001 Smith expanded his network of co-practitioners to George Mason University where he is now practicing experimental economics. But by this time the experimental method had become an established field of economics.¹⁰¹

The resistance to experimentation forced Smith to come up with persuasive arguments that could justify the use of the method in economics. Thus, along with his experimental practice Smith started to adduce arguments that could justify the experimental method of economics.

Why do we do what we do in the ways that we do it? None of my research outside of experimental economics required me to come to terms with this question. When you do mainstream research everyone takes for granted your right to do it. In confronting these issues I was led to write about them... This sort of writing seems inevitable if you want to convey some appreciation for the role of experiments in increasing economic understanding. It is not enough to just do experiments, and to report them, because they are not part of the culture of your audience. (Smith 1990, p. 5)

In the previous sections it was argued that the first published reports of market experiments not only reported experimental results but also attempted to provide arguments that justified the use of the experimental method in economics. It was also shown how the conceptual understanding of market experiments and the experimental method evolved with Smith's practice. However, a full-blown account of the use of the experimental method in economics has not yet been supplied.

Given that the experimental method had been nonexistent in economics, its justification had to have recourse to items of scientific culture from other fields of research. The mechanism design theory lent itself as an obvious resource in that it could easily be presented as the theoretical counterpart of experimental economics.

The justification of the use of the experimental method in economics relied heavily on the field of research that Stanley Reiter (1977) dubbed the (new)² welfare economics and, in particular, on Leonid Hurwicz's 'Informational Efficiency of Resource Allocation Mechanisms' (1960).¹⁰² It provided a formal framework for the

study of resource allocation mechanisms which became "the unknown of the problem, rather than a datum" (Hurwicz 1960, p. 28). The similarity between the two research programs is evident. The (new)² welfare economics aimed at providing a theoretical scheme for the comparative appraisal of alternative allocation mechanisms in terms of their various performance characteristics (e.g. optimality properties, informational efficiency, etc.). In addition, it was intended as a scheme for the design of resource allocation mechanisms on the basis of pre-defined criteria. Thus, whereas the (new)² welfare economics provides a theoretical framework for the design and evaluation of resource allocation mechanisms, the experimental method provides an experimental framework to the same effect.¹⁰³

Even though the (new)² welfare economics gave experimental economics an identifiable and acceptable reference upon which the objectives of experimental economics could be presented and justified, the adoption of this framework to the construction of a theory of experimental economics was not immediate.

Initially, the strategy was to attract the interest of mechanism design theorists in the use of the experimental method to test market mechanisms. This is very clear in Smith's co-authored article with Plott (1978) where both research programs are presented as being concerned with "the analysis of information and organization in decentralized price adjustment processes". The experimental method, in particular, was presented as "a 'proving ground' for theories" (p. 133) and "an empirical challenge for those interested in price adjustment processes" (p. 147). The appeal to mechanism design theorists continues in (Smith 1980b) as Smith proposes to "develop a foundation for the study of resource allocation mechanisms" (p. 345) in an article with the expressive title 'Relevance of Laboratory Experiments to Testing Resource Allocation Theory'.¹⁰⁴

The contribution of (new)² welfare economics to help justifying the function and role of experiments in economics is finally made explicit by Louis Wilde (1981). Wilde not only explicitly presented experimental economics as the experimental counterpart of the new approach to welfare economics, but he also adopted the conceptual framework of the latter to economic experiments. It was Wilde who first noted that "a laboratory experiment in economics attempts to create and study a small-scale microeconomic environment... [whose] purpose is to uncover systematic

¹⁰² The (New)² Welfare Economics distinguishes from the New Welfare Economics in that the latter's "main accomplishment . . . was to derive conditions characterizing (Pareto) optimal production and exchange, and to show that in classical economies (those in which preference and production possibilities have suitable convexity and continuity properties) the equilibria of the competitive mechanism precisely meet the conditions characterizing Pareto optimal production and exchange". The (New)² Welfare Economics concerns the evaluation of alternative systems where "the allocations which are the outcome of the system are just one of the important aspects of its functioning, but others are also important, among them its administrative feasibility, the costs of operating the system itself, the extent to which private incentives are incompatible with the system" (Reiter 1977, p. 226).

¹⁰³ It is no coincidence that the parallel research programs emerged at about the same time. Welfare economists were also very much inspired by the research at RAND Corporation, in particular, after being disappointed with the potential of mathematical statistics for guiding policy advice. Indeed, the founders of both 'new' versions of Welfare Economics (Kenneth Arrow, Jacob Marschak, Tjalling Koopmans, Leonid Hurwicz) were all research associates at the Cowles Commission which had direct contact with the RAND Corporation. Mirowski talks about a 'conversion experience' operating at the Cowles Commission that resulted in "looking seriously at game theory and linear programming, reorienting its efforts toward a general theory of organization patterned on the computer" (2002 p. 264).

¹⁰¹ For a biography of Vernon Smith see <http://www.gmu.edu/departments/economics/facultybios/smith.html> and <http://nobelprize.org/economics/laureates/2002/smith-autobio.html> and other links therein.

relationships between individual preferences, institutional parameters, and outcomes” (p. 139). Thus, Wilde had a very important role in bringing together the relevant items of scientific culture to construct an articulated framework that justified the use of experiments in economics. From Wilde (1981) it was only a small step for Smith to round up the extant contributions and put forward a theory of laboratory experiments as he did in his “Microeconomic Systems as an Experimental Science”. To conclude, the early arguments Smith put forward in defence of experimental economics relied on the persuasive force of the similarity between the theoretical research program and the experimental method. The upshot was, or so it seems, that experimenters are simply doing what theoreticians do by different means.

However compelling, Smith’s strategy did not attract the interest of the theoreticians. The resistance of the economists to the experimental method might be explained by the fact that the goals of the experimental program represented a paradigmatic shift from the bulk of the work done at that time. Specifically, it implied a rejection of the standard assumption of unbounded rationality and the background framework of perfect competition.

Experimental Economics seemed to be more appealing to the younger generation of economists inspired by the new field of game theory. Those who were receptive to the game theoretical paradigm were also those who were more receptive to the experimental method. The subsequent development of experimental economics shows just that. The view that game theory was the major driving force of experimental economics is now commonly accepted in the field:

Game theory brought to economics a kind of theory that lent itself to experimental investigation, and in some cases demanded it. The reason is that it seeks to provide precise models of both individual behavior (in von Neumann-Morgenstern utility functions) and of economic environments. This concern with the “rules of the game”, the institutions and mechanisms by which transactions were made, together with precise assumptions about the behavior of individuals and the information available to them, gave rise to theories that could be tested in the laboratory. (Roth 1995a, p.18)¹⁰⁵

In addition to game theory, the experimental work in economics has also been stimulated by the problem-situations the earlier sets of experiments give rise to. That is, like experimentation in the natural sciences, experimental economics also possesses a ‘life of its own’. The use of the experimental microeconomic systems to explaining worldly phenomena or to providing guidance for policy-making has been far more limited, however.

It should also be noted that Smith had recourse to non-economic resources to argue for the uses of the experimental method. Of particular interest is the invocation

of ‘non-experimental’ sciences, such as astronomy and meteorology, which, like economics, can make the results of micro-scale experiments (e.g. mass motion and of the thermodynamics properties of gases) to bear on macro-scale phenomena (e.g. stars and the climate) (1982, p. 936). From philosophy of science Smith adopts the Kaplan taxonomy to present the functional and the methodological uses of laboratory experiments (op. cit., p. 940, see section 8.2 below). But the most effective, perhaps, was the falsificationist rhetoric that pervades the seminal paper, and which was quite popular among economists (cf. Blaug 1992, p. xiii). This point deserves special attention in the next chapter.

5.9 CONCLUSION

The analysis of Vernon Smith’s trajectory is illustrative in showing that the establishment of a new field of research is the outcome of a long process of knowledge production that involves a wide community of researchers. Experimenters had to become full experts on the new method before they were able to build the methodological foundations for it. In order to do that, they had also to have recourse to items of scientific culture from other fields of research as well as from other disciplines.

Even though the chapter focused on Smith’s practice, it is evident that the foundation of the experimental field in economics is a collective achievement. The methodological foundations of the method, as put forth by Vernon Smith in his 1982 seminal paper, are the result of coherent achievements that articulate his experience and that of others. But, more importantly, experimental economics owes much to the social process of criticism and to the scepticism encountered. It was this scepticism that motivated Smith to produce robust results that could face up the criticism of the sceptics and ultimately to put forward an articulated theory of experimentation in economics.

Given that the experimental method was deemed inadequate in economics, the justification of the first experimental results was accompanied by an extensive array of arguments that justified both the results and the method. Even though the construction and invocation of relations of coherence were necessary to establishing the field of research, they were not sufficient. The resistance was too acute. The growth and development of the field owed instead to the potential relations of coherence that the younger generation could foresee for future practice. Notwithstanding the resistance of the community of economists to the method, the wider socioeconomic environment was favourable. It welcomed its use for ‘institutional engineering’ that could not be undertaken by other means.

The foregoing analysis showed that economic experiments have a ‘life of their own’, too. Initially, the market experiment was designed and implemented to test a theoretical hypothesis, but soon it gave rise to a whole new range of problem-situations that could be explored by experimental means. In the course of experimental practice, the designs, the procedures, the stylised facts of experimental economics accumulated and gradually formed the scientific culture of the field of research. Experimental economics is now a well-established sub-discipline of

¹⁰⁴ The same strategy is followed in the field of public choice theory. Plott argues that similarly to (new) welfare economics, which aims at designing ‘resource allocation mechanisms’ according to pre-defined performance criteria, public choice theory engages in ‘institutional engineering’ that aims at constructing “‘new’ or ‘synthetic’ institutions which have prespecified performance characteristics” that “may or may not resemble any existing institution” (1979 p. 139).

¹⁰⁵ See also Friedman and Sunder (1994, p. 122).

economics. It has a well-delineated area of research which can make use of a set of well-established items of experimental culture. However, the relative autonomy of experimental economics has contributed to an excessively defensive attitude towards external criticisms and contributions. As will become apparent in the following chapters, the methodological foundations of the field are not yet well grounded nor are the epistemic functions of economic experiments clearly understood.

The application of the social epistemology of experiment to the market experiments revealed that this category of experiments tend to be characterised by a low level of human agency, a high level of stringency and a high level of technological applicability. The market experiment constitutes a stringent experimental system that produces unambiguous results because experimenters are able to control subjects' behaviour. It was also shown that the high level of control exercised over human behaviour is the main factor responsible for the technological applicability of the market experiments in real world environments. From this it follows the broader conclusion that the higher the degree of 'human agency', the lower the level of stringency of economic experiments and therefore the lower the level of technological applicability. However, a more active participation of experimental subjects in knowledge production carries a higher potential for generating knowledge about human behaviour. This topic will be postponed until chapter 7. The foundations of the experimental method as put forward by Smith need further scrutiny. This is the topic to which I turn to in the next chapter.

CHAPTER 6

THE EPISTEMOLOGICAL ARGUMENTS OF EXPERIMENTAL ECONOMICS

6.1 INTRODUCTION

Chapter 5 presented a stylised view of experimental economics. It was introduced as the field of research that creates and controls microeconomic systems in the laboratory either to investigate theoretical microeconomies or to learn some aspect about real world ones. This chapter scrutinises in detail the characteristics of the experimental method of economics. In particular, it analyses the arguments experimenters put forward to justify: 1) the possibility of creating and controlling microeconomic systems in the laboratory, 2) the adequacy of experimental microeconomic systems to theoretical testing and development, and 3) the adequacy of laboratory microeconomic systems to providing knowledge about real world microeconomies. I will then show that these arguments were biased by an ill-formed view of science and experimentation. The social epistemology of experiment framed in chapter 4 will then be brought to bear on this discussion and thereby refine, revise or reformulate the arguments that justify the adequacy of the experimental method for its various purposes.

The chapter is organised as follows. Section 2 describes in detail how and for which purposes economists create and control microeconomic systems in the laboratory. Section 3 puts forward the *social epistemological argument* that justifies the epistemic value of experimental microeconomies. This argument is then brought to bear on the various functions that experiments may serve. Section 4 analyses the 'simplicity' and the 'artificiality' criticisms levelled against experimental economics. Section 5 addresses the transferability of experimental results to concrete real world situations. Section 6 scrutinises the social dimension of the experimental practice of economics. Section 7 provides a sketchy portrait of economic experiments. Section 8 concludes the chapter.

6.2 HOW ECONOMISTS CONTROL EXPERIMENTAL MICROECONOMIES

Control is the hallmark of scientific experimentation and experimental economics is no different in this respect. Therefore, the possibility of experimenting in economics calls for the scrutiny of the way by which economists create and control microeconomic systems in the laboratory.

FOR WHAT PURPOSES DO EXPERIMENTERS WANT TO CREATE MICROECONOMIES?

As we have seen in the previous chapter, the goal of creating a manageable microeconomic system in the laboratory is to study the four-fold relationship *microeconomic environment – microeconomic institution – individual behaviour – aggregate outcome* that had long been central to economic analysis but unobservable and undetectable by non-experimental means. The major difficulty here concerned the measurement of microeconomic environmental variables, in particular, individual preferences, individual knowledge and technological endowments. The experimental method allowed, for the first time, the study of the relation between the attributes of economic agents, the set of rules of communication and exchange whereby agents interact with one another, the actions they take under those rules, and the performance of the microeconomic system. As Smith put it:

Among the observable elements of an economy are (i) the list of agents, (ii) the list of physical commodities and resources, (iii) the physical commodity and resource endowments of individual agents, (iv) the language and property right characteristics of institutions, and (v) outcomes. What is not observable are (vi) preference orderings, (vii) technological (knowledge, human capital) endowments, and (viii) agent message behavior... These last items are not observable because they are not only private but to a degree *unrecorded*. Willingness to buy (preferences) and willingness to produce (technology and preferences) can at best only be inferred from agent point actions in the message space. Often we cannot even observe point messages, for example, we may know allocations and prices, but not all bids. In any case, we cannot observe the message behavior *functions* because we cannot observe (and vary) preferences. (1982, p. 928, emphasis in original)

As Smith leaves patent, field data, i.e. data obtained from uncontrolled and naturally occurring processes,¹⁰⁶ provide information about economic agents, institutions and aggregate outcomes, but this information does not illuminate on certain issues relevant to economics. Consequently, and in contrast to experimental data that are purposively generated to bear on particular problem-situations, field data constitute only indirect evidence for the study of the four-fold relationship of interest. For instance, field data does not allow determining whether the rules that are specified in the institution in conjunction with the behaviour of the agents, yields a choice of messages which constitute, say, a Nash equilibrium whose outcomes are Pareto optima.¹⁰⁷ This is because this assessment requires information about individual preferences and production possibility sets that is not obtainable by non-experimental means.¹⁰⁸ At best, crude proxies can be obtained from the available data. But these cannot aspire to be more than very unreliable estimates of the relevant variables. A substantial amount of conceptual and statistical work must

be done to convert the field data into estimated values for the desired variables. This comparative advantage of economic experiments in relation to other methods is well captured in Smith's assertion:

We can measure certain preference and/or technology parameters... and the effect of certain institutional rules... Furthermore, within the specifications of the model... we can test particular hypotheses... Rarely are we able to obtain a test of the model specification. Hence, an econometric model provides a mapping from specifications into conclusions about preferences, technology, and institutions. Insofar as the conclusions are sensitive to the specifications, we are left with scientific propositions that are open-ended with respect to the environment, institutions, and agent behavior. Furthermore, since parameter identifiability and the properties of estimators depend upon model specification, the particular model chosen inevitably must be influenced partly by the technical requirements of the methodology and not only of the scientific objectives of the exercise. (1982, p. 929)

The problem is that the work required to convert non-experimental data into useful information about the problem at hand is highly sensitive to scientists' prior beliefs. In the face of ambiguous results, there is a high likelihood of improving the model's fit with "reasonable expectations". As a result, "the whole process becomes an exercise in fitting a particular belief system to field data by manipulating model specifications and perhaps estimation methods" (op. cit., p. 929, n. 8).¹⁰⁹ Of course, as Smith acknowledges, experimentation is not exempt from the same practice, but the high level of social robustness does limit it. The high level of susceptibility of experimental results to public scrutiny contributes to the identification of experimenters' priors and their subsequent test (see section 5.6). To put it in another way, because the level of 'human agency' of experimental systems is higher than that of econometric models, it allows the more direct measurement of variables pertaining to human behaviour. In addition, the high level of 'sociality' of the experimental process of production of knowledge attenuates the effect of the economists' prior beliefs in the production of experimental data.¹¹⁰

¹⁰⁸ Surveys are not an option here. A major difficulty concerns the translation of economics concepts into intelligible and meaningful questions. The difficulties associated with the use of survey data for microeconomic analysis is well illustrated by the marginalist controversy surrounding the test of the assumptions of "perfect competition" or "perfect monopoly", on which R. E. Lester (1946, 1947) and Fritz Machlup (1946, 1947) were the main contenders.

¹⁰⁹ Indeed, the permeability of econometric results to scientists' priors was at the centre of a huge crisis in econometrics in the 1980s (see Mayer, 1980; Leamer, 1983). This crisis has been attenuated by the development of more sophisticated procedures to account for the above-mentioned problem and by the promotion of a more methodologically self-conscious practice. At any rate, the comparative advantage of the experimental method to control and observe preferences remains.

¹¹⁰ This not to mention field data's susceptibility to non-scientific goals given that this kind of data is often collected by non-scientific agencies for non-scientific purposes. As Smith put it, "when things appear not to turn out as expected the quality of data is more likely to be questioned than the relevance and quality of abstract reasoning". In contrast, the experimental process of knowledge production increases the credibility of scientific data. It "brought to the economist direct responsibility for an important source of scientific data generated by controlled processes that can be replicated by other experimentalists" (1987, p. 242).

¹⁰⁶ See Friedman and Sunder (1994, p. 3-7) for a more complete taxonomy of different sources of data.

¹⁰⁷ A Nash equilibrium is a situation where each agent chooses a strategy that is the best (utility-maximising) response to the other agents' strategies and a Pareto optimum is a situation where no one can increase his utility without decreasing that of others.

HOW DO EXPERIMENTERS CONTROL MICROECONOMIC SYSTEMS?

Being the hallmark of the experimental method, it is not surprising that control was the first methodological issue addressed by the experimenters (Smith 1976a, 1976b, 1980b, 1982; Plott 1979, 1982; Wilde 1981). The ‘institution’ of the microeconomic system is easily controlled by design. It consists of the rules of communication and exchange that the experimenter defines and enforces during the course of the experiment. The major methodological difficulty experimenters face concerns the control of the ‘environment’ of the microeconomic system, namely, the preferences of the experimental subjects. According to Induced Value Theory (1976a) control over preferences can be exercised via a reward structure that induces prescribed monetary value on experimental variables. To that effect the payoff function must take as arguments the experiment’s variables and subjects must be given the property right over the monetary outcomes of the actions taken during the experiment. The use of money to induce value on abstract experimental outcomes is very ingenious in that it presupposes the same principles prevalent in real economic systems. As Smith notes, in the economic system agents have also property rights over “intangible property on which value is induced by specifying the rights of the holder to claim money or goods” (1982, p. 931, n. 11).

In addition, the reward structure must satisfy “the set of sufficient conditions for a valid controlled microeconomic experiment”: nonsatiation, saliency, dominance and privacy (cf. section 5.2). In the double auction experiment, for instance, value is induced on the commodities transacted by a reward structure that pays the buyers the redemption values of the goods bought and pays the sellers the market price of the units sold, and by property rights rules that confer subjects the right to claim the amounts earned with these transactions (obtained after deducing the prices buyers paid and the cost of production of the goods sold). The precepts of nonsatiation and saliency are satisfied insofar as the amounts subjects earn by participating in the experiment are important to them and these amounts are related with the experimental outcome. That is, the rewards earned vary positively with the gains from trade.

Because there might be nonmonetary subjective costs and values inherent to subjects’ participation in the experiment, the reward structure may be insufficient to guarantee effective control. In particular, there may be the “subjective cost of transacting, that is, the cost of thinking, calculating, and acting” which may offset the value induced by the reward structure. In addition, “individuals may not be autonomous own-reward maximizers” for “interpersonal utility considerations may upset the achievement of well-defined induced valuations” (op. cit, p. 933-5). In order to exercise control over these factors the conditions of dominance and privacy must also be verified. These conditions ensure that the amounts earned are significant to the subjects so that they compensate for the costs inherent to the process of making and executing decisions.¹¹¹ By omitting the overall pattern of the reward structure, intersubjective considerations are avoided.

In sum, in a microeconomic experiment subjects’ preferences are controlled by way of a reward structure that prizes successful decisions made in a context of relative anonymity. In this way self-interested reward-maximising motivations are induced on experimental subjects. The goal is to ensure that the experimental situation is relevant to the subjects in the intended way so that the experimental microeconomies are intelligible systems. Only then can experimenters understand how individual preferences and institutional rules translate into individual behaviour and how this, in turn, affects the performance of the microeconomy. As Smith put it:

When great care is used in an experiment to make induced value be the primary source of motivation, it is *not* for the purpose of making sure that subjects have self-interested motivations; it is for the purpose that we *know* what were the preference patterns of the subjects in the experiment. (1982, p. 933, n. 13, emphasis in original)

Even though the inducement of self-interest and income maximising motivations is not a methodological prescription of experimental economics, the importance of the postulates of self-interest and income-maximisation in economic theory, and the need of controlling subjects’ preferences so as to create transparent experimental systems, has rendered these *de facto* the major motivational patterns of experimental economics.

The need of controlling subjects’ motivations substantially reduces the agency of the experimental subjects in economics. In chapter 2, in the context of experimental practice in the natural sciences, it was argued that the direct participation of the object under scrutiny in knowledge production conveys epistemic value to experimental results. It carries the potential to ‘resist’ scientists’ expectations and therefore to prevent experimental products from being exclusively determined by scientists’ material and conceptual actions. In economics, the tight control exercised over subjects’ preferences via the reward structure and the additional control exercised over subjects’ actions via the microeconomic institution substantially reduce the ‘resistance’ of ‘human agency’ (recall Smith’s double auction experiment). But it does not follow from this that ‘human agency’ does not resist or that the epistemic value of economic experiments calls for a different explanation. It only means that the role of ‘human agency’ in experimental economics needs closer inspection.

Strict adherence to the methodological prescriptions of experimental economics does not ensure that individuals are self-interested income-maximisers or that they succeed in taking the course of action that best suits their interests and/or that some desired social goal is reached. For one thing, subjects’ motivations are multiple and not always effectively controlled by the experimenter. For another, the cognitive limitations of the individuals and the complexity of the decision-making process may prevent them from perceiving and pursuing the best course of action. And even when perceived and pursued, the best courses of action may produce undesirable collective outcomes. It is precisely this potential mismatch between individual preferences and individual actions (e.g. the winner’s curse), or between individual actions and aggregate outcomes (e.g. collusion), where the relevance of microeconomic experiments lies. The goal of experimenting with microeconomic systems is precisely to study the factors that affect

¹¹¹ This points to the importance, in any experiment, of designing simple and transparent experimental contexts so that subjects can easily understand and execute what is required from them.

the alignment between individual preferences, individual actions, and some desired social goal. In addition, an economic experiment may also have the goal of controlling human behaviour for certain purposes.

We are now in condition of summarising the role ‘human agency’ plays in experimental economics:

The participation of experimental subjects is the main source of epistemic value of an economic experiment. It is the major factor that causes ‘resistances’ to economists’ expectations and thereby contributes to the growth of knowledge. Insofar as experimenters control human behaviour by inducing self-interest and income-maximising behaviour, the ‘human agency’ of economic experiments allows the identification of the circumstances that affect the manifestation of this behaviour. From this it follows that the epistemic goal of a microeconomic experiment is to learn under what circumstances human behaviour deviates from self-interested maximising action or to learn how to control it for socially desired ends.¹¹²

THE EXPERIMENTAL PRINCIPLES AND PROCEDURES OF ECONOMICS¹¹³

As explained above, a controlled microeconomy is a transparent system that elicits behaviour that can be interpreted in the light of preferences induced by the reward structure and the microeconomic institution that organises subjects’ interaction. However, the design of a reward structure that fulfils the precepts of experimental economics does not suffice to warrant that experimental subjects have understood the experimental situation as intended by the experimenters. A larger set of principles and procedures must be followed to create transparent or stringent experimental systems. Four general principles that guide the practice of experimental economics may be distinguished. These principles are simplicity, neutrality, abstraction and anonymity.

Simplicity is a central principle of experimental economics. The design of simple tasks is crucial to guarantee that subjects understand the experimental task in the intended way and that they can perform it within the available time. Misunderstandings on the part of the subjects render the results of experiments unintelligible and therefore they jeopardise the whole enterprise. Simple tasks also minimise the costs associated with the effort of executing experimental tasks and therefore prevent behaviour caused by insufficient attention or fatigue. An important

procedure in this respect consists in conceding subjects enough time to become acquainted with the experimental situation. Subjects should have sufficient time for reading the instructions, for practicing trial tests, and for clarifying misunderstandings. After the conduction of the experiment, questionnaires can be used to appraise subjects’ level of comprehension.

The increasing use of computers in experimental economics has contributed significantly to the simplification of the experimental tasks. The use of computers allows presenting experimental tasks in a more intelligible way and reduces time for executing them. Moreover, it facilitates implementing and enforcing the rules of the experiment because forbidden actions may be excluded by the computer program. This possibility allows subjects identifying mistakes and thereby reducing error. The use of computer also allows the immediate record of experimental data, which reduces error with data collection and treatment. The simplicity of experimental tasks presents also the practical advantage of reducing the costs associated with running the experiment in that it may reduce the monetary rewards to compensate subjects’ effort with the execution of experimental tasks.

The principles *neutrality* and *abstraction* control the interference of individuals’ subjective perceptions about the context of interaction. This is achieved by avoiding the use of terms loaded with suggestive connotations (e.g. maximisation, competition, collusion) and references to concrete real world situations. In addition to minimising the effect of subjects’ prejudices and misconceptions, neutrality also allows the more effective concealment of the experiment’s goals. Consequently, it renders more difficult for subjects’ to act in conformity, or defiance for that matter, to what might be perceived as the ‘right action’.

The *prohibition of deception* is an important rule of experimental economics that aims controlling subjects’ perceptions about the experimental context. The goal is to develop among the subject population a reputation for honesty so that subjects may be led to trust the stated purposes of the experiment. In this way the effect of subjects’ conjectures about what the purposes of the experiment might be is minimised. Other practices are followed to the same end. For instance, economists follow closely the script of the experiment to avoid the effect of unintended factors and, whenever feasible, they reread the relevant part of the instructions to answer to subjects’ questions. Embarrassing the subjects is also avoided throughout the experiment. To this end, subjects are paid in private so that they do not know each other’s payoffs.

Anonymity is other important principle of experimental economics. It aims at controlling the effect of intersubjective considerations, i.e. subjects caring about the experiment’s outcome to other subjects. In addition to the implementation of the precept privacy, anonymity is promoted by avoiding as far as possible contact between subjects. Thus, the structure of the experiment must take into account subjects’ interaction outside the laboratory prior to the experiment, during breaks, and after the experiment is finished. The use of computers in experimental economics has an important contribution here, too. It reduces both the contact between the experimenter and the subjects as well as among subjects. This is an important feature to guarantee the reproduction of experimental results in that the effects of subject pools and of experimenters are significantly reduced.

¹¹² Even though the analysis follows Smith’s conception of a microeconomic experimental system that more directly bears upon the epistemic role of market experiments, the same analysis applies to individual decision-making and game theory experiments. Considering the role of ‘human agency’ in economics experiments, and simplifying a bit, one could say that individual-choice experiments investigate the factors that affect individual rationality in contexts of uncertainty or risk, and game experiments investigate the factors that prevent individuals from taking dominant strategies, or the factors that influence the attainment of cooperative and noncooperative solutions to strategic problems. In the next chapters I will illustrate and explain how ‘human agency’ manifests in these two categories of experiments.

¹¹³ This section follows the textbook presentation of the experimental procedures of economics, namely Davis and Holt (1993, Chs. 1 and 9) and Friedman and Sunder (1994). It should be noted that the principles and procedures to be described apply to the stylised view of experimental economics presented in chapter 5. Therefore, they may not apply equally to all the experimental work done in the field. As will become apparent in subsequent chapters, their relative importance depends on the purpose of the experiment.

There are, however, other factors that may not affect subjects understanding of the situation but may have an impact on the experiment's results. These factors raise problems of confounding effects by being correlated with the treatment variables. Factors that may be expected to be strongly correlated with particular variables may be controlled directly by *blocking*, that is, by including them as treatment variables.¹¹⁴ Other background factors may be indirectly controlled by *randomisation*, that is by assorting at random background factors to treatment variables.¹¹⁵

Finally, experimental economics also make use of statistical techniques that continue the experimental work after the experiment had been run. Many experiments in economics generate data that provide straightforward answers to the problem-situation at hand. However, and depending on the problem-situation, in many experiments substantial statistical tests and data analysis must be undertaken to investigate whether and to what extent the data support the hypothesis under test (or does not disprove it).¹¹⁶ That is, statistical tests and data analysis investigates the extent to which there is a good fit between the predicted and the observed outcomes, whether the difference between data generated by different treatments is significant, the testing of explanatory hypotheses, and so forth.¹¹⁷

To conclude, it is clear by now that the principles and the procedures of experimental economics aim at creating fairly aseptic systems shielded from the contamination of the outside world. However, it should also be clear that the isolation of experimental systems from the outside world can only be achieved to some degree. After all the agency of human beings, both in the laboratory and in the real world, is where the interest of conducting experiments with them lies.

CONTROL AND THE SOCIALITY OF EXPERIMENTAL ECONOMICS

The way by which control is to be achieved in experimental economics is now well ingrained in the field. The conceptualisation of a microeconomics experiment, the 'set of sufficient conditions for a valid microeconomics experiment', and their general principles and procedures are all part of the established culture of experimental economics. This can be easily ascertained by the textbook depiction of the field of research (Davis and Holt 1993; Friedman and Sunder 1994; Hey 1991) or by the field's methodological discussion (e.g. Loomes 1991, Starmer 1999b, Guala 2001, 2005b) and, most expressively, by the experimental disputes that often revolve around

¹¹⁴ For example, if experience is considered such a background factor, its effect may be isolated by running a section with experienced subjects and other with inexperienced subjects.

¹¹⁵ Subjects order of arrival to the experiment, subjects' relationship with one another or the experimenter, the order of treatments are all examples of factors that can be controlled by randomisation.

¹¹⁶ Statistical tests allow the making of negative claims only. That is, while they can support the rejection of hypotheses they do not allow accepting them when they fail to be rejected.

¹¹⁷ For the purposes of the present chapter it suffices to mention that statistical tests and the techniques of data analysis constitute an important set of procedures of experimental economics. The choice of the tests and techniques to the analysis of experimental data depends on the design of the experiment, namely the type of data it generates (binary, discrete, continuous), the nature of the sample (matched or independent observations) and the structure of the hypothesis (single parameter, multiple, or an entire distribution). See Davis and Holt (1993, Ch. 9) for a survey on the tests and techniques used for particular designs.

the conformity to these standards (see Ch. 9). This means that the processes of knowledge production, evaluation and dissemination have been successful in enforcing what may be labelled as the 'method of experimental economics'.¹¹⁸

No doubt the high degree of exposure of the process of knowledge production to public scrutiny facilitates this enforcement. This exposure, as mentioned earlier, is the result of the fairly complete reporting of economic experiments in scientific journals. Moreover, it is standard practice of experimental economics to render relevant information available upon request as a way to enhance the credibility of experimental data (Davis and Holt 1993, p. 22). In addition to the instructions, experimenters may also have access to information regarding the recruitment of subjects, the implementation of trial tests, the composition of the subject pool and their experience level, the procedures used for matching subjects to roles, the location, dates and duration of experimental sessions, the use of laboratory assistants, computers and other special devices, and the experiments' data.

The standardised nature of economic experiments not only enhances the transparency of the experimental systems but it also favours the replicability and the robustness of experimental results.¹¹⁹ Indeed, the adoption and adaptation of previous instructions and experimental procedures in the design of new experiments is a pervasive practice in experimental economics (to be developed in section 6.6).

In conclusion, the creation and control of microeconomic systems is realised by well-established methodological procedures that include:

1. The definition of a salient reward structure capable of inducing the desired set of preferences on the experimental subjects.
2. The creation of fairly neutral and abstract experimental contexts to avoid the interference of uncontrolled factors.
3. Guaranteeing anonymity among subjects to avoid the effect of interpersonal considerations.
4. The use of stationary repetition whenever adequate to ensure that subjects understand and acquire experience with the experimental situation.
5. The prohibition of deception to enhance the efficacy of the experimental instructions (in present and future practice).
6. The convention of providing fairly complete reports of experimental procedures.

6.3 THE SOCIAL EPISTEMOLOGICAL ARGUMENT: EXPERIMENT VERSUS THEORY

It was argued that the methodological procedures of experimental economics aim at creating transparent and intelligible experimental systems. However, these do not

¹¹⁸ Recall that this thesis focuses on outcome-based economic experiments (cf. footnote 63, p. 97) or the stylised view of a microeconomic experiment.

¹¹⁹ As defined in Ch. 4, replicability refers to the reproduction of the same experiment and robustness to the reproduction of experimental results with different experimental systems.

suffice to accept the relevance of the experimental systems for intended purposes. Further arguments must be provided to justify the adequacy of economic experiments for each of the specific purposes they may serve. This section scrutinises one of the functions of economics experiments – the empirical testing of theories.

As we have seen in chapter 5, Smith's first experiment (1962) aimed at testing competitive price theory. At that time no argument was presented to justify the adequacy of the market experiment to that effect. The experimental market was simply presented as an operational model of a competitive market, as depicted by CPT, that could be used to check whether the experimental market converged to equilibrium. Because convergence was observed, Smith concluded that the theory predicts well therein. The relevance of the exercise remained unexplained. In 1982, Smith presented the microeconomic experiment as a 'falsifying' tool to become another twenty years later, a more modest instrument for theoretical construction and development. In the process, more elaborated arguments were adduced and presented to justify the relation between experiment and theory.

Falsificationism (Popper 1959, 1965) has been and still is a very popular methodology among economists.¹²⁰ It is therefore not surprising that experimental economists also resorted to falsificationist arguments to justify the use of experiments. According to falsificationism, good scientific practice consists in the proposal of 'bold conjectures' and in their submission to 'severe testing'. A 'bold conjecture' is a hypothesis that makes a low-probability prediction about the world and a 'severe test' is a deliberate and tenacious attempt to produce evidence that falsifies the theory. If the test generates negative evidence, the theory is 'refuted' and, as a result, it should be discarded and replaced by a new one. If the hypothesis survives the test, the theory is 'corroborated' by experience which means that the theory has resisted attempts at falsification (rather than being confirmed).

Smith does not elaborate on the use of economic experiments as falsifying tools. Smith simply alludes to the falsifying power of economic experiments. He asserts that experimental microeconomies are adequate testing devices because they carry the potential for falsifying theories. However, underlying his argumentation seems to be the idea that the falsificationist potential derives from the 'reality' of the microeconomies, or their higher level of 'richness', 'complexity', and 'concreteness' in comparison to theories:

Precepts 1-4 [nonsatiation, saliency, dominance, privacy] permit us to study laboratory microeconomic environments in which *real* economic agents exchange *real* messages through *real* property right institutions that yield outcomes redeemable in *real* money.

Insofar as we are only interested in testing hypotheses derived from theories, we are done, that is Precepts 1-4 are sufficient to provide rigorous controlled tests of our ability as economists to model elementary behavior. Microeconomic theory abstracts from a rich variety of human

activities which are postulated not to be of relevance to human economic behavior. The experimental laboratory, precisely because *it uses reward-motivated individuals drawn from the population of economic agents in the socioeconomic system*, consists of a far *richer* and more *complex* set of circumstances than is parameterized in our theories. Since the *abstractions* of the laboratory are orders of magnitude smaller than those of economic theory, there can be no question that *the laboratory provides ample possibilities for falsifying any theory we might wish to test*. (1982, p. 935-6, emphasis added)

Smith does not explain how and to what extent the 'reality' of microeconomic systems (or its increased levels of richness, complexity, and concreteness) enhances the falsifying conditions of the experimental system. But from what was argued in the previous chapter it is not difficult to see that the falsifying possibilities of experimental systems derive from 'human agency'. That is, the possibility of having experimental subjects behaving differently from the behaviour postulated by economic theory, or the possibility of microeconomic institutions affecting human behaviour in unexpected ways. In other words, the falsifying conditions of an economic experiment depend on the 'resistance' of the human world to economists' expectations.

The mere participation of human beings does not suffice to justify the adequacy of an experiment as a falsifying tool. To be a good falsifying tool, the experiment must constitute a genuine attempt at falsifying the theory under test. That is, the design of the experiment not only must allow, but it must also trigger falsifying behaviour. Only then can an experiment be considered a 'severe test'. To be sure, a microeconomic experiment is not an adequate test if it constrains human agency in such a way that renders virtually impossible the manifestation falsifying behaviour. As I argued in the last chapter this was the case of Smith's early market experiment. The observations were not very compelling evidence because the 'experimental procedure' significantly induced subjects' behaviour and thereby the experimental results. The participation of economically motivated subjects in a well-defined experimental situation is not sufficient to guarantee 'ample possibilities for falsifying any theory'. Nor is it the implementation of a transparent design that establishes a clear relation between the hypothesis under test and the observations consistent with it. In addition to creating a context where the experimental results can be easily interpreted, the way by which these results come about must be the result of the agency of the experimental subjects and/or that of the microeconomic institution. A more adequate argument to justify the use of economics experiments to test theories is instead:

An economic experiment is a severe test of hypothesis H if and only if H implies observation O, and there is a very low probability of observing O in the event that H is false. Insofar as the observations generated by economic experiments are produced by human agency, there is room for observing individual behaviour and aggregate outcomes that conflict with O. However, the following condition must be verified: the experiment must not constrain human behaviour in such a way that observing O is virtually ruled out by design.

Notwithstanding the apparent triviality of this observation, this is an issue that has not been noted by the experimenters or the students of economic experiments. As will

¹²⁰ Falsificationism is to-date the most influential methodological proposal in economics. Wade Hands interprets the appeal of falsificationism to economists in terms of its simplicity and adoptability insofar as it provided a "set of easily implemented methodological rules for the proper conduct of scientific inquiry" (Hands 2001, p. 276). Hands also argues that Popper's interest in the social sciences, Popper's work in political philosophy, and Popper's personal and professional connections at LSE might have also contributed to the importation of falsificationist ideas to economics.

become apparent below, and in the next section, experimenters have assumed at the outset the adequacy of an experiment for empirical testing. They have even pointed out that a good test gives the theory the ‘best shot’. The students of experimental economics have in turn overlooked the specificity of experimental economics. That is, the fact that experimentation in economics consists of intervening in the human world.

Falsificationism (or any methodology based on empirical testing, for that matter) is afflicted by the well-known problem of underdetermination, or the Duhem-Quine thesis, which undermines the epistemic power of empirical testing.¹²¹ The thesis put in evidence the fact that the relation between an observational statement and a theoretical hypothesis is not of a deductive kind which undermined the possibility of falsifying theories on the basis of refuting observational statements.

The Duhem-Quine (DQ hereafter) thesis asserts that the test of an empirical hypothesis is always a conjoint test of a target hypothesis (e.g. a competitive market tends to price-quantity equilibrium) together with a variety of auxiliary hypotheses (e.g. nonsatiation, salience, dominance and privacy are satisfied in the laboratory). As a result, experimental data are inconclusive.¹²² Negative data to the hypothesis under test do not refute it. The data may constitute instead evidence for the falsity of the auxiliary assumptions. Conversely, a positive result does not provide evidence that the target hypothesis has been corroborated. To put it in another way, evidence is as fallible as theory is, and therefore a clash between theory and evidence does not have sufficient disproving force. This is specially the case if the theory is well established. Consequently, when experimental data conflict with ingrained and expected beliefs, scientists question empirical results instead. The fragility of empirical testing implied by the DQ thesis was not ignored by the experimental economists. In subsequent methodological work, Smith shows that he is fully aware of it:

When the data are consistent with the predictions of a theory, it is sometimes said that the results are not interesting because they merely confirm what economists already knew... which seems to suggest that “truly” authoritative theory cannot be doubted seriously. When the data are inconsistent with the predictions of theory it is not uncommon to assert that there must be “something” wrong with the experiments... Similarly, questions about experimental procedure are more likely to be raised when the results appear to disconfirm accepted theory than when they appear to confirm such theory. (1989, p.167-8)

From a falsifying instrument, the microeconomic experiment became a tool for theoretical ‘extension’ with increased ‘empirical content’ within the realm of a given research program. The Lakatosian methodology of scientific research programs is now the model of good scientific practice (Lakatos 1970):

¹²¹ The problem of underdetermination or the Duhem-Quine thesis is named after the physicist Pierre Duhem (1954 [1906]) and the philosopher Willard Orman Quine (1951).

¹²² Indeed, economists’ first reaction to refuting evidence is to blame the experiment’s design and the failure of the precepts of experimental economics: payoffs are too small, subjects do not have experience, subjects do not have the opportunity to learn, instructions are ambiguous, etc. (these and other illustrations are provided in subsequent chapters).

In any confrontation between theory and observation the theory may work or fail to work. When the theory works it becomes believable in proportion to its predictive “miracle”, instead of only respectable in proportion to its internal elegance or its association with authority. But when it works, you lean mightily upon the theory with more challenging “boundary” experiments designed to uncover the edges of validity of the theory where certainty gives way to uncertainty and thereby lays the basis for extensions in the theory that increase its empirical content... When the theory fails to work in initial tests, the research program is essentially the same. This is because all theories can be expected to be more or less improvable, and statistical tests of theories, whether the results are initially “falsifying” or not, are simply the means to motivate extensions in theory. Better theory that narrows the distance between theory and observation is always welcome. (Smith 1989, p. 152)

Not only could the more tenable Methodology of Scientific Research Programs account for the actual practice of economists, who did not really strive for rejecting their own theories, but it could also demonstrate that experimental economics has been part of progressive research programs and that experimenters do just what they should. Whatever the result of the experiment, the ultimate goal is to increase the theory’s empirical content. This can be done either by pushing the edge of the theory’s validity when it resists falsification or by modifying the theory in the light of refuting evidence. Notwithstanding the increasing methodological awareness, the justification of the use of a microeconomic experiment for theoretical development had yet to be provided. Thus far the DQ problem had only been acknowledged and the goals of experimental practice adjusted accordingly.

The backbone of an answer appears in Smith (2002) where a broader view of the experimental process of knowledge production is presented. Not only does Smith reinforce the view that scientists do not abide by the falsificationist stricture and that “what scientists do is better than what they say about what they do”, but he also adds that what they do is guided by “[t]heir *natural instincts* and *lively professional interactions* [that] lead them perpetually to design *new experiments that examine the right questions*” (op. cit., p. 91, emphasis added). Underlying this statement is the presupposition that the problems raised by the DQ problem can be solved by the conduction of more experiments and that the relevance of these depends on the making of the ‘right questions’. This, in turn, depends on experimenters’ ‘instincts’ and ‘lively interactions’ within the community of experimental economists. More explicitly, Smith states:

[E]xperimentalists design new experiments with the intention of confronting the issues in the controversy, and in the conflicting views that have arisen in interpreting the previous results. This leads to new experimental knowledge of how results are influenced, or not, by changes in procedures, context, instructions and control protocols. The new knowledge may include new techniques that have application to areas other than the initiating circumstance. This process is driven by the D-Q problem, but practitioners need have no knowledge of the philosophy of science literature to take the right next local steps in the laboratory. Myopia here is not a handicap... The bottom line is that good-enough solutions emerge to the baffling infinity of possibilities, as new measuring systems emerge, experimental tool kits are updated, and understanding is sharpened. (Smith 2002, p. 103-4)

To conclude, in addition to the relative superiority of the experimental method in improving the transparency of the experimental systems, and hence in generating superior inferences, the experimental method allows examining remaining controversy by experimental means. The collective nature of the experimental process of knowledge production is fundamental in this respect (cf. Ch. 3). It brings about alternative interpretations that may be explored and discriminated by conducting new experiments. We are now in condition of synthesising the epistemic value of microeconomic experiments in the form of the *social epistemological argument of experimental economics*:

The epistemic value of a microeconomic experiment derives from the 'agency' of experimental subjects and from the 'structure' of the context wherein they act and interact. These create ample possibilities for the generation of results that 'resist' experimenters' expectations. In addition, the openness of microeconomic experiments to the scrutiny of the wider scientific community contributes to identifying and solving pertinent questions to the problem-situation at hand. Because the levels of 'agency', 'structure' and 'sociality' are variable, the epistemic worth of the economic experiment depends on the extent to which the process of knowledge production created opportunities for both 'human' and 'social' resistances.

Notwithstanding the fact that DQ problem is an inescapable feature of science and that experimenters, consciously or not, do know how to deal with it in actual practice, there remains some anxiety about the absence of an objective arbiter between theory and evidence and the resulting controversy in experimental economics. To illustrate this, it is worthwhile to quote in length Chris Starmer's views on the controversy surrounding the preference reversal phenomenon:

[T]here are some areas of research where it is harder to see what the, sometimes quite extensive, volume of research has accomplished. Consider, for instance, research in relation to the preference reversal phenomenon... Preference reversal looks like quite a serious issue: it not only violates EU [expected utility theory], but seems to pose quite a challenge for economic theory in general... 'It suggests that no optimization principles of any sort lie behind even the simplest of human choices'... After a good deal of research on this phenomenon, what do we know? For instance: is preference reversal a real phenomenon? Surely we can agree that it does occur in specific experimental contexts. But, is it robust to increases in payoffs; does it occur in economically meaningful contexts; how should it be explained? Each of these questions remains highly controversial, or so it seems to me. Moreover, despite the very large volume of effort devoted to empirical investigation of this phenomenon there seems to have been very little effort developing or applying economic accounts of the phenomenon. Under these circumstances, an outsider observer could perhaps be forgiven for reaching the conclusion that the sustained empirical effort involved here has had little impact upon economic thinking more broadly. (1999b, p. 24)¹²³

The significance of DQ problem is given by a misguided view of science that sees experiments as neutral arbiters in the settlement of disputes and dispute resolution as a unanimous and final endpoint to the controversy. However, the fact that no

such arbiter exists does not mean that disputes cannot reach a consensual resolution. Neither does consensus require unanimity. In fact, the recognition of DQ problem renders disagreement not only acceptable but also desirable and unanimity a far more worrying epistemic state.

Moreover, as will become apparent from the analysis of experiments presented in subsequent chapters, the major causes of disagreement in experimental economics do not concern the interpretation of experimental results but instead the inferences experimenters want to derive from them. As noted above, the high level of stringency of microeconomic systems not only renders economic experiments very transparent systems, so that the source of ambiguity can be identified and examined with more experiments, but it also renders experimental results very dependent on the design of the experiment (e.g. the double auction market experiment and market equilibrium) which significantly constrains the inferences afforded by the experimental method. Under these conditions, controversy is above all the result of weakly grounded inferences that do not take into account the specificity of the data generating process, or a sign that the work undertaken has not been sufficient to clarify the points of contention. In other words, the state of affairs that Starmer complains about says more about the community of economists which has not addressed the issues Starmer would have liked, or has made inferences that are not supported by the experiments realised. Of course, it may also be the case that the research program has reached a dead end. Scientific happy endings are as a mythical as crucial experiments. At any rate, the epistemic superiority of experimental data relative to other empirical data remains. Auxiliary hypotheses can be subjected to the test and the range of admissible interpretations narrowed down.

To summarise, the epistemic value of experimental economics stems from the possibility of creating stringent experimental microeconomies and from the potentially high level of social robustness of the process of knowledge production. The high degree of stringency renders experimental microeconomies fairly transparent processes of generation of data that reduce the interpretative range of experimental results. The high level of social robustness, in turn, contributes to identifying and discriminating among the conflicting interpretations. The epistemic significance of experimental results depends on the level of 'human agency' of experimental systems. It depends on extent to which the experimental results were generated by the actions of the experimental subjects, placed in a particular socioeconomic context, rather than the actions of the experimenters.

6.4. THE SOCIAL EPISTEMOLOGICAL ARGUMENT: THE 'SIMPLICITY' AND 'ARTIFICIALITY' DEBATES

The community of economists have doubted, and many still do, that experiments can provide any relevant knowledge about economic behaviour or economic institutions. These doubts are often associated with what are here labelled as the *simplicity* and *artificiality* criticisms. While the former targets the simplicity of the laboratory and its inability in capturing the complexity of 'real world' environments,

¹²³ The preference reversal program is the object of Ch. 9.

the latter denies the possibility of recreating natural occurring processes in the artificial environment of the laboratory.

Before addressing the criticisms straight on, it should be noted that the debate has been rather asymmetrical in that the critiques levelled against the experimental method of economics have not been translated into fully articulated and published arguments. The criticisms appear, only implicitly, in experimenters' written responses to them.

In general terms, the *simplicity criticism* points to an insurmountable distance between the experimental microeconomy and the 'real world' microeconomy the experiment purportedly refers to. This is how Plott and Smith depicted its overall tone:

Criticism of experimental methods... often seems to arise from the assumption that what is important about an experiment is that it be "real worldlike" in some subjective sense, informally conjectured by the investigator. Critics frequently have in mind some sort of "simulation" of a complicated naturally occurring situation and look to the experiment to provide forecasts about how events will unfold in the future (perhaps under various scenarios). (Plott and Smith 1979, p. 168)

The distance between the experimental situation and the 'real world' has been interpreted differently, however. Two alternative interpretations are clearly identified. One argues that the experimental situation can never capture the complexity of real-world situations. The experimental microeconomy is a simple and neutral context of interaction in which subjects guided by self-interest and reward-maximising motivations make fairly abstract decisions. In real-world contexts, in contrast, human motivations are heterogeneous and dependent on the specificity of the context of interaction. Indeed, the attention experimenters dedicate to the design of the experiment demonstrates just that. The major methodological concern of experimental economists is precisely to control human motivations, which are multiple and dependent on the context. Given the obvious parallelism of this criticism with the heterodox critiques of the so-called orthodox economic theories, this is dubbed the '*heterodox*' *simplicity criticism*.¹²⁴ The other criticism is that the experimental context does not even create the necessary conditions to generate self-interested motivations and maximising behaviour. Because this criticism is voiced by economists in the face of refuting evidence to economic theory that presupposes self-interested maximising behaviour, it is dubbed the '*orthodox*' *simplicity criticism*.¹²⁵

The *artificiality criticism* presents a more radical view of the experimental situation. The problem of experimental economics is not the inadequacy of experimental subjects

to represent the behaviour of economic agents. It is rather the fact that the behaviour generated is substantially different from the behaviour occurring in a 'real world' situation. The point is that human behaviour is highly sensitive to context and therefore it will differ in different environments such as the laboratory and the real world environment. Taking up the criticism raised against experimental psychology, Nikos Siakantaris reiterates that the "laboratory is not a socially neutral context, but is itself an institution with its own formal or informal, explicit or tacit rules" and therefore "the laboratory is itself a special kind of society" and one distorted by "the idea of experimenting with humans" (2000, p. 274-5).

The artificiality criticism hinges on an important distinction between the natural and the human sciences. In the human sciences, scientific experimentation involves human beings who are conscious individuals whose behaviour depends on how they perceive the situation they are in and how they decide to behave in it. Consequently, the artificial context wherein subjects interact inevitably affects subjects' perceptions and actions.

To put it differently, the *simplicity criticism* targets the 'material' quality of economic experiments, i.e. their human and socioeconomic 'richness' or 'complexity', to use Smith's terms. The *artificiality criticism* targets the relevance of the 'material' component of economic experiments. That is, while the former argues that the agency of experimental subjects cannot be manifested in its full extension in the laboratory, but may be useful for more modest purposes (simple problem-situations), the latter is arguing that the agency experimental subjects manifest in the laboratory is not meaningful.

RESPONSES

The very practice of experimental economics denies the significance of the problems associated with the artificiality criticism. The experimental endeavour presupposes that the experimental method is capable of producing meaningful knowledge. Underlying this presupposition is a belief in the capability of creating and controlling microeconomic systems in the laboratory and, in particular, in controlling factors that may cause unintended and undesirable perceptions on experimental subjects, which are at the centre of the artificiality criticism. Given that the method of economics has been concerned with how best to achieve control over these factors, as explained above, it is not surprising that the experimenters have built confidence that the behaviour they observe can be interpreted in terms of relevant variables to the experiment. This confidence is expressed in experimenters' responses to the simplicity criticism:

[E]conomies created in the laboratories might be very *simple* relative to those found in *nature*, but they are just as *real*. *Real* people motivated by *real* money make *real* decisions, *real* mistakes and suffer *real* frustrations and delights because of their *real* talents and *real* limitations. Simplicity should not be confused with reality. Since the laboratory economies are *real*, the general principles and models that exist in the literature should be expected to apply with the same force to these laboratory economies as to those economies found in the field. (Plott 1991, p. 905, emphasis added)

¹²⁴ 'Orthodox' economic theories are often charged of being too abstract (rather than applied to concrete real world situations), of adopting narrow or false assumptions about individual preferences and human cognitive capabilities (e.g. the assumption that individuals are solely guided by self-interest and that they possess unbounded rationality), of neglecting the influence of the institutional and the social setting on individual and aggregate behaviour, among others. This debate has obvious connections with economists' diverging views on what is the purpose and the adequate level of realism of economic theories. For economy of space, this topic is to be addressed only in passing below. For an overview of this debate in economics see Mäki (1992, 1994, 1996, 1998, 2002).

¹²⁵ Graham Loomes provides a very eloquent depiction of this criticism by the orthodox skeptic (see Appendix 3 at the end of this chapter).

Experimenters dismiss the artificiality criticism head on by claiming that the experimental microeconomic system is as ‘real’ as a natural socioeconomic system because ‘real people’ make ‘real profits’ within the context of ‘real rules’ and therefore the principles of economics that apply to the natural contexts also apply to the laboratory. This assertion was already implicit in the falsificationist account of experiment, but here it is used to argue for the similarity between the ‘artificial’ or the ‘simple’ experimental context and the ‘natural’ referent in the real world. However, this does not constitute an adequate response. The distinction between the natural and the social sciences that the artificiality criticism stresses is not taken into account. The ‘real people’ in economics experiments may still behave differently in the experimental context. But neither does the invocation of that distinction, as will be argued below, justify the irrelevance of experimental observations. Before coming to that, the answers to the simplicity criticism are scrutinised.

EXPERIMENT AS A TESTING TOOL OF THEORIES

The response to the simplicity criticism introduces a distinction between the two main goals of a microeconomic experiment that attempts to neutralise the criticism for theory-testing while acknowledging its relevance for policy-making (Plott 1978, 1982, 1987, 1990; Plott and Smith 1979; Smith 1980, 1982; Wilde 1981). The argument is that as long as an experiment aims at testing a theory, or at discriminating between theories, it does not need to reproduce, and no presumption need be made about its connection with the more complex ‘real-world’ context. It only needs to include the parameters relevant to the theory or theories being tested. The comparative advantage of an experiment is precisely to study interesting relations confounded in nature. In the words of Smith:

What is important about an experiment is that it be relevant to its purpose, not that it be realistic in the sense that it be “real-world-like” in some subjective sense. Indeed, the best experiment is the crucial experiment whose outcome clearly distinguishes between competing theories. But the conditions of the crucial experiment may rarely, if ever, occur in nature. (1980b, p. 350)

Hence, the relevance of the experimental setting relies on the possibility of creating an experimental situation that can be brought to bear on interesting theoretical problems. This possibility in turn lies on the ability of the experimenters in controlling subjects’ motivations so that human behaviour can be interpreted as intended by the design of the experiment. In particular, it depends on the ability of the experimenters in inducing the motivations postulated in economic theory (e.g. self-interest motivation) and of being assured that they succeed in doing so. In fact, it is this possibility that supports the experimentalists’ assertion that ‘the microeconomic system is a real microeconomy where real agents engage in the resolution of real economic problems’.

EXPERIMENTS SUBSERVIENT TO THEORIES

The resistance of the economics community led experimenters to provide additional arguments to justify the adequacy of microeconomic experiments for theory-testing. A further argument directs the simplicity criticism to the theory and to the critics:

Experiments are sometimes criticized for not being “realistic”, i.e. parallelism is questioned. There are two appropriate responses to this criticism: First, if the purpose of the experiment is to test a theory, are the elements of alleged unrealism in the experiment parameters of the theory? If not, then the criticism must be directed to the theory as much as to the experiment. Laboratory experiments are normally as ‘rich’ as the theories they test. Second, are there field data to support criticism, i.e. data suggesting that there may be differences between laboratory and field behavior. If not, then the criticism is pure speculation; if so, then it is important to parameterize the theory to include the behavior in question. (Smith 1980b, p. 350)

The argument here is the following. Because an experiment is an instantiation of the theory in the laboratory, it necessarily possesses a lower degree of abstraction and simplicity (as explained above) and therefore it is at least as ‘realistic’ as a theory is.¹²⁶ Therefore, charges of ‘unrealisticness’ compromise the theory rather than the experiment. And if empirically grounded, these charges put in evidence a failure of the theory which should then be revised.

Of course, it is good scientific practice to substantiate criticisms. Evidence and arguments should be provided in their support. But the strategy of directing the burden of criticism to theory is not warranted. First, this is an overly defensive attitude that may prevent dialogue and thereby the continuity of the knowledge production process. As already noted, experimenters must be open to criticism which is the basis of new testable hypotheses that may be discriminated by designing and conducting new experiments. Moreover, experimenters must be able to justify their interest in carrying out particular tests to particular theories as much as they must be able to justify the adequacy of experimental tests for specific purposes. An experimenter is not a mere technician who waits for the theoreticians’ testable hypotheses. The practice of experimental economics simply denies it.

Second, experiments possess a degree of autonomy from theory which may render irrelevant theoretical controversy. On the one hand, any theory, even a widely accepted one, may give rise to a multitude of experimental problem-situations. Theories do not come with testing instructions that specify what is to be tested and what a satisfactory test is. On the other hand, experiments are autonomous processes for generating data whose relevance can be independently justified. Experimenters should instead focus on the identification of interesting problem-situations that can

¹²⁶ Smith is not clear about what a ‘realistic’ theory amounts to. However, he seems to endorse the view that the representations of economic theories can and should capture the relevant elements of the domains they purportedly apply to (e.g. the ‘fundamental structure of the economy’ or the ‘the way the world works’, cf. Mäki, 1998, p. 312-3). If some of these are left out or false they should be included or revised.

be adequately solved by experimental means. If the experiment design is adequate for the purpose at hand then the relevance of the experiment is justified.

Third, the justification of a given experiment should not depend on the endorsement of a particular philosophy of theory. Indeed, Smith seems to endorse a realist conception of theories which is not necessary to justify the use of experiments. Different philosophies of theory are consistent with the experimental method (more on this below).

EXPERIMENTS AS TESTS OF GENERAL THEORIES

Plott introduces a qualification to the use of experiments for theory-testing. The argument is that simple laboratory microeconomies are good test devices of *general* theories or models. General theories depict economies ‘found in the wild’ by representing their structure and by involving “basic principles intended to have applicability independent of time and location except to the extent to which time and location have an effect on... preferences, institutions, information, and feasible sets”. The simplified experimental microeconomies provide good testing conditions for these general models because “[g]eneral models, such as those applied to the very complicated economies found in the wild, must apply to simple special cases” (1991, p. 905). Accordingly, if the theory fails in these special cases then it also fails in the complex economies found in the wild.

Once models, as opposed to economies, became the focus of research the simplicity of an experiment and perhaps even the absence of features of more complicated economies became an asset. The experiment should be judged by the lessons it teaches about theory and not by its similarity with what nature might happen to have created. (Plott 1991, p. 906)

This is a skilful way to answer the simplicity criticism. Rather than a handicap of the experimental method, the simplified nature of experimental systems is a methodological asset insofar as it reinforces the epistemic weight of negative evidence to a theory of general application. If the theory does not predict well in simple environments then it should be discarded or modified. But Plott’s views incur in the same problems noted above. First, it renders the justification of economic experiments contingent on negative results. Second, it makes the justification of experiments contingent on a particular conception of economic theory. Specifically, it restricts the range of application of the experimental method to the test of theories of a certain kind – general theories – leaving unjustified the use of experimentation for other purposes such as theoretical development and the test of non-general theories.¹²⁷ At best, it presupposes that the domain of application of economic theories is explicit so that it can be ascertained whether or not a given theory is expected to apply to the laboratory. However, this is not the case of most economic theories which do not have explicit domains (Lipsey and Crystal 1995, p. 31). At any rate, the interest of experimental economics is precisely to test the edges of

validity of theories, models and conjectural hypotheses. Experimental practice would be a very limited enterprise if it only tested theories in the domains where it is already expected to apply (see also Cubitt 2005).

THE SOCIAL EPISTEMOLOGICAL RESPONSE

The debate around the simplicity criticism seems to constitute the experimental counterpart of the debate around the ‘realism’ or ‘realisticness’ of economic theories. In fact, the charges that used to be levelled against the abstract and formal theories of economics are now directed against the abstract and aseptic nature experimental contexts, and their simplified assumptions or induced behaviours. Oddly enough, however, these charges may come from both fronts. While economists with a preference for more descriptively accurate theories may perceive the experimental situation as an excessively simplified context, economists who may accept high levels of abstraction in economic theory, may consider that an economics experiment is not an adequate test because it does not provide subjects with enough conditions to behave like the economic agents postulated in economic theories (cf. Loomes’ caricature in appendix 3, see also footnote 124). This may help explaining why experimenters were so keen in repeating “experimental microeconomies are *real* economies” where “*real* people pursue *real* profits within the context of *real* rules”.

But the invocation of the ‘reality’ of experimental microeconomies does not suffice to justify the adequacy of experiments as testing devices. The interpretation of ‘real’ as a context where experimental subjects possess the motivations postulated by economic theory and where they engage in decision-making problems analogous to real-world problems does not do the trick. As the analysis of the market experiment demonstrated, an experiment that merely implements the behaviour postulated by the theory is not an adequate testing device. However, if by ‘reality’ one means ‘human agency’, then the experiment has epistemic value. As the *social epistemological argument of experiment* leaves patent, the epistemic value of economic experiments derives from the agency of experimental subjects because it prevents experimental knowledge from being the exclusive result of experimenters’ actions.

The notion of ‘human agency’ dissolves the problems that the notions of ‘reality’ or ‘realisticness’ bring about. Not only is the ‘reality’ of experimental subjects insufficient to justify the epistemic worth of economics experiments, but ‘realisticness’, understood as a complete and faithful representation of a real economic situation, is not very illuminating either. In the first place, the conceptualisation of ‘reality’ as conveyed by the three-fold condition “real people - real profits - real rules” is not very elucidating if it is meant as the creation of a context wherein ‘real rules’ by way of ‘real profits’ induce on ‘real people’ the behaviour postulated by the theory while substantially constraining the behaviour that would lead to its refutation. Secondly, there will always be significant differences between the experimental system and its ‘real-world’ counterparts (e.g. experience, stakes, and so forth), which cannot be easily settled (e.g. how much experience is realistic enough?). Moreover, if the experiment attempts to mirror a real-world situation in all its detail and complexity it will not yield conclusive results. In short, the results of the experiment can always be dismissed either

¹²⁷ This problem is intensified by the absence of an argument that justifies preference over general theories (cf. Hausman 1992 and Guala 2005b, Ch. 7).

because experiments are ‘unrealistic’ in some important sense, or too realistic and therefore uninformative. The relevant question to ask then is: How adequate is the experimental design for the purpose at hand? And the answer to this question hinges on the assessment of the level of participation of ‘human agency’ in knowledge production.¹²⁸

The ‘human agency’ of economic experiments justifies the use of laboratory experiments to generate knowledge about human behaviour regardless of the goal of the experiments, and regardless of whether the theory is realistic or instrumental, local or general, or whether a parallel situation exists in the ‘real-world’. The relevance of economic experiments depends on the ability to create and control experimental systems where variables can be manipulated in interesting ways and experimental subjects can understand the experimental task as intended. Insofar as observed behaviour manifests the agency of experimental subjects, the significance of experimental results is established.

As the analysis of experimental economics showed, an important characteristic of the precepts and the procedures of experimental economics is that they aim to produce very transparent contexts (which are the grounds of the simplicity and artificial criticisms) that substantially narrow down the range of interpretation of the experimental systems. If some ambiguity prevails, it can be eradicated by further experimenting.¹²⁹ The fact that the experimental situation does not resemble any known real world situation does not undermine the value of an experiment to learn about human behaviour or about the effect of institutional frameworks on human behaviour and social outcomes. Of course, if the experiment is designed to have some bearing on some real world situation it has to resemble it in important respects. In this case, the experimental situation must include the relevant and essential items of the real world situation and represent them in the laboratory. (This issue is addressed in next section)

The ‘realism’ debate in experimental economics shows that the adequacy of critical charges is as relevant as that of the responses to them. The debate about the pertinence of economic experiments has been guided by an ill-informed view of experiment based on a theoretically oriented conception of science. Both criticisms and responses were not very illuminating. Methodological discussion in economics has focused on the representational attributes of economic theories. Economists have questioned the abstract (as opposed to concrete applications), the narrow (excluding many), the neglect (excluding most relevant) or the falsehood of economic explanatory factors. When transposing this debate to experimentation,

economists have implicitly presupposed that an experiment must as well be an adequate representation of some existing real-world situation. However, the view that a microeconomic system must represent a real-world socioeconomic context does not take into account the role of a central component of the experimental system: the participation of human subjects in socioeconomic problem-solving. This participation introduces an important distinction between an economic theory and an economic experiment: whereas an economic experiment is a real socioeconomic situation, an economic theory cannot aspire to be more than a representation of a real-world situation. How far and the extent to which this situation may have some bearing on other situations is a legitimate research-question, but a subsequent one. The relevant epistemic question of experimental economics is rather: *How far and to what extent did the experimental system create sufficient conditions for the agency of experimental subjects?*

THE ARTIFICIALITY CRITICISM REVISITED

The focus on the ‘reality’ of economics experiments and the relative confidence on experiments for theory testing has been the cause of perplexing results. While experimenters seem comfortable with the use of experiments for theory-related goals, they are uneasy with the use of experiments for other purposes. This tension is very clear in the next passage by Arthur Schram when referring to the use of experiments to establishing ‘stylised facts’:

[I]t appears to me that the artificiality of these experiments [bargaining experiments] is very high. Therefore, it is not clear what the value is of this documentation of empirical regularities. If the aim is to make any claims about other regarding preferences in the world at large, the high artificiality appears to render the experimental results useless. If, in contrast, the aim is to document robust, causal laboratory effects to confront theorists with... artificiality is less of a problem. This does have the danger of theorists and experimentalists creating their own world, however...

For gathering empirical regularities that aim at telling... something about behavior outside of the laboratory, it is important that the artificiality of the situation is considered. Much more than in case of testing theories, the data must be relevant for the situation one is interested in. (2005, p. 233)

Notwithstanding the legitimate concern that inattention to real world phenomena may substantially limit the potential of experimental economics, the economists’ confidence that experiments are partial means of assessing and improving the applicability of theories to the real world and their scepticism about the use of experiments to provide knowledge about the economic world are perplexing. This perplexity suggests that the empirical quality of economic experiments, present in theory testing, is somewhat absent when the experiment is used for other purposes. But if the epistemic value of experiments is conveyed by the participation of human beings in the process of knowledge production, as argued

¹²⁸ This issue is common to theoretical economics. Given that any theory is inescapably a partial depiction of reality insofar as it focuses on a slice of reality and it is by necessity an abstract and idealised representation of it (for reasons of tractability), it will always be unrealistic in some sense. Thus, the relevant question to ask in theoretical economics is the adequacy of the theory for the purpose at hand. As Mäki put it, “since all theories contain unrealistic assumptions, the real issue is –and should be construed as– one about the substance of those theories and assumptions, namely what they exclude as irrelevant or inessential and what they include as relevant or essential, and what they say about the included items” (1998, p. 312-3, emphasis in original).

¹²⁹ In the remainder of the thesis examples will be provided that show how and the extent to which alternative interpretations have been discriminated against by the conduction of more experiments.

here, this value is independent of the experiment's goal. The fact that experimental results are produced by the decisions and actions of human beings implies that they generate knowledge about human behaviour that may be applicable to the real world. As argued above, how far and the extent to which this knowledge applies to some concrete real world situation is a relevant but independent research question.

Of course, the requirement of transparency limits the sort of knowledge that can be produced by economic experiments (see section 6.2). Even though many interesting phenomena cannot be implemented in the laboratory, the information provided by experiments may nonetheless improve our understanding about them. This is explicitly recognised by Nicholas Bardsley (2005), who has thus far provided the most in-depth analysis of the 'artificiality criticism'. Bardsley excludes 'relational' phenomena, i.e. phenomena that "depend jointly on certain relationships obtaining between people and on people's perceptions that these same relational criteria are satisfied" (p. 241). The problem, as he sees it, is that "a laboratory experiment cannot implement them unchanged without deception, since subjects are aware of the additional, generative activity of the experimenter. The subject-experimenter relationship, which is common knowledge in a non-deceptive experiment, implies that the laboratory generates its own SPP [social psychological phenomenon]" (ibid). From this it follows that the artificiality of experiments is not solvable by design in some spheres of human life. More specifically:

The strict identity of laboratory and target variables, from which natural scientific experiments derive their demonstrative force, does not hold. There are two aspects of laboratory experiments which allegedly undermine this identity. Both arise from the experimenter-subject relationship. That gives rise to the non-identity of certain laboratory relationships and institutions with their real counterparts (where an experiment attempts to implement them) and the peculiar theory-ladenness of experimental stimuli (a general concern). (op. cit., p. 245)

For instance, the tax evasion experiment, where subjects are asked to report the amount of money received by the experimenter and on the basis of which they pay an experimental tax, cannot provide information about tax evasion as it occurs in the real world. None of the decisions taken in the experiment can be interpreted as instances of tax evasion because "a tax is revenue collected by a government, which has an authority relationship over a group of citizens. Moreover, for an action to qualify as tax evasion, it must be seen as such by the agent". To be sure, a "certain episode of behaviour (typically, writing down a figure on a tax form) qualifies as tax evasion only if it is done with the intention of paying less tax than the system requires. There is therefore a reference to taxation in the agent's intentions in real tax evasion that is absent from the actions in the experiment" (op. cit, p. 242).

Bardsley concludes that these results are more adequately taken as "discoveries about how people behave in laboratory public good games with a probability of being punished for non-contribution" because "[p]eople might recognise a civic or legal duty to pay taxes whilst not recognising a duty to be honest to experimenters in labs, or indeed vice versa" (ibid). But Bardsley does not exclude the possibility that the results of this experiment may bear on tax evasion, the point is that the

experiment alone does not support that inference. Thus, Bardsley is not arguing that experiments are always artificial and therefore no interesting knowledge can be derived from them. He is making the milder claim that some real world phenomena cannot be reproduced in the laboratory. The social events that can be implemented, in turn, pertain to events that "can, in normal life, be organised amongst strangers without a particular institution's involvement" (op. cit., p. 247).

Bardsley's account is very illustrative in showing that an economic experiment is a social context which must be taken into account when deriving conclusions from it, or to put it in another way, experimental results are context-dependent. The social epistemology of experiment highlights that the participation of experimental subjects, whose actions depend on their subjective perceptions and interpretations, bestows epistemic value to those results whether or not there is a real-world situation to which these results may be brought to bear. In conclusion, the focus on the role of human subjects in economic experiments dissolves the relevance of the goals of experiments to the epistemological analysis of experiments, as well as that of the relation between the experiment and some referent in the real world. The bottom line is that economic experiments allow acquiring knowledge under very simple and artificial environments because they allow observing how individuals behave in them.¹³⁰

6.5 THE PARALLELISM PRECEPT: EXPERIMENT VERSUS THE REAL WORLD

Besides the empirical testing of theories, experiments can serve other research goals such as the study and interpretation of field data, the study of existing institutions, or the design and implementation of new ones. The adequacy of the experimental method to any of these purposes is supported by the parallelism precept that justifies the application of experimental results to nonlaboratory contexts. In particular, it asserts that "(p)ropositions about the behavior of individuals and the performance of institutions that have been tested in laboratory microeconomies apply also to nonlaboratory microeconomies where similar *ceteris paribus* conditions hold" (Smith 1982, p. 936). In the context of market experiments, for instance, parallelism concerns the possibility of generalising the propositions about experimental market institutions to nonlaboratorial markets:

[I]f institutions make a difference, it is because the rules make a difference, and if the rules make a difference, it is because incentives make a difference... whatever the context of the particular microeconomy... parallelism says that the incentive effects of different bidding rules are *qualitatively* the same; if rule A produces lower bids than rule B in one market, it will do so in other markets. (Smith 1982, p. 936-7, emphasis added)

The claim to generality is first and foremost a qualitative claim that pertains to the relation between the environment, the institution and the performance of the system.

¹³⁰ This topic is explored in chapter 8.

For more quantitative results, parallelism would require a tighter correspondence between the laboratory and the nonlaboratory environments (p. 937). In any case, parallelism is an empirical matter that demands further empirical work.

In contrast to the use of experiments for theory testing, experimenters are very cautious about the extrapolation of experimental results from the laboratory. Plott is very explicit in clarifying that the full understanding of the ‘simple’ microeconomic system does not allow the immediate extrapolation of the results yielded to the more ‘complex’ cases:

Behavior in very complex environments may follow different laws than those which govern behavior in relatively simple situations. This is an excellent reason for being careful in any attempt to extrapolate behavior from a laboratory to a complex industry... It is the most difficult task that any researcher faces. Experiments are simply an additional source of data and experience that one adds to other sources in making judgments about how the world works. (1982, p. 1522-3)

In Plott’s later work, it is even suggested that knowledge about the ‘naturally occurring processes’ is more adequately generated by other methods, and that the experimental method is simply instrumental to that end:

I do not believe that experimental methods will replace field research. Economies found in the wild can only be understood by studying them in the wild. Field research is absolutely critical to such an understanding. However, the theories and models used in field research necessarily incorporate many judgments about assumptions, parameters and behavioral principles. The simple cases that can be studied in the laboratory can provide the data against which the importance of such judgments can be assessed. (1991, p. 918)

Notwithstanding the fact that the generalisation of experimental results to non-laboratory contexts is the most challenging issue experimenters are confronted with, it is also the methodological issue of experimental economics that has been most studied. The work of the philosopher Francesco Guala (1998, 1999b, 2001, 2002a, 2002b, 2003, 2005a) represents such contribution. Guala proposes to examine parallelism as a problem of *external validity*, to be opposed to the less complex inferential problem of *internal validity*.¹³¹ More formally:

Internal validity is achieved when the structure and behavior of a laboratory system (its main causal factors, the ways they interact, and the phenomena they bring about) have been properly understood by the experimenter. For example: the result of an experiment E is internally valid if the experimenter attributes the production of an effect B to a factor (or set of factors) A, and A really is the (or a) cause of B in E. Furthermore, it is *externally valid* if A causes B not only in E, but also in a set of other circumstances of interest, F, G, H, etc. (2003, p. 1198, emphasis in original)

However, Guala notes, the precept of parallelism is an external validity claim of a very particular kind – it is a local kind of inference:

Parallelism inferences are thus local in character, from particular circumstances to particular circumstances. External validity (from inside the laboratory to any unspecified outside system) does not capture such an intuition. Speaking of parallelism instead of external validity can perhaps help us to mark an important distinction. To achieve parallelism is best characterized as a weak form of external validity limited to some typical variations in the background conditions. It is to establish that a certain experimental result can be “exported” from a set of (experimental) systems E to a set of (“real-world”) systems W. It is a *local* inference, since economists do not have to believe that the relationships they subscribe to must hold across *all* kinds of situations. (1999b, p. 569-70, emphasis in original)

The exportability of experimental results to real world environments represents an ‘analogical inference’ that requires the scrutiny of the similarity between the experimental system and the target system. According to Guala, this scrutiny consists, firstly, in the comparative analysis of the observable features of the experiment and the target system in order to assess whether the similarity can be accounted for by the same causal factors. Secondly, the experimenters must eliminate alternative explanations to the observed similarity.

The idea, roughly, is that we need to create (or select) circumstances in which it is really unlikely to observe certain data, unless the external validity hypothesis is true. In this case, the data is the correspondence between observed features of the target and observed features of the experimental system; the external validity hypothesis is that the data belong to similar causal mechanisms. Now, the probability of observing such a correspondence (were the hypothesis false) is low if we have eliminated alternative reasons why such a correspondence might occur, other than the causal similarity between the two systems... if you want to generalize from A to B, you better make sure that A and B are as similar as possible. (2003, p. 1202)

Parallelism is most likely to be verified in cases where the target system can be accurately reproduced in the laboratory or, alternatively, the experimental system can be reproduced in the field. Guala says:

“Exporting the lab”... is just one route to external validity – the safest one perhaps, but not the *only* one. Its viability depends on how much we are allowed to intervene and shape reality to fit the experimental prototypes... a more general methodology [is available], that can be applied (with varying degrees of reliability) also when the laboratory cannot be exported, but must be adapted to the target at hand. (2003, p. 1204)

On a similar vein, Nikos Siakantaris considers that parallelism is most likely to apply to systems that exist in relative isolation in the real-world:

A necessary condition for parallelism to hold is that the experimental situation set up in the laboratory is analogous to the real-world one. But that an analogous situation exists, in some broad sense, is not enough for parallelism. It is also necessary that the analogous situation exists in conditions of relative isolation outside the laboratory, which in turn implies that it be relatively stable and identifiable... The implication is that parallelism is more likely to hold for real-world

¹³¹ This terminology is imported from experimental psychology.

situations that exhibit considerable stability and are not influenced significantly by some other factors. (2000, p. 273)

The limitation of parallelistic claims to well-circumscribed real-world systems is in fact supported by extant practice. The known cases of parallelism refer to auction markets that are well-defined and stable target systems that exist in conditions of relative isolation outside the laboratory. But not only does parallelism apply to phenomena that exist in relative isolation, it also applies to phenomena that can be faithfully reproduced in the laboratory (e.g. winners' curse) or, once produced in the laboratory, can be reproduced in the field under similar conditions (e.g. FCC auction). Thus, parallelistic claims are restricted to claims about human behaviour in very rigid environments or about institutions that are very effective in controlling human behavior for intended purposes.¹³² In any case, the extrapolation of results from the lab to real world situations is a complex empirical matter of a rather limited scope.

The limited range of parallelism is a manifestation of what Siakantaris calls the 'experimental economics trade-off': "The better experimental economists do their job in controlling variables, the more they are threatened by a lack of parallelism and hence of the usefulness of their project" (2000, p. 273-4). But this way of presenting parallelism is misleading. As we have seen, control is also what allows the transferability of the results of economic experiments to the real world. The exportability of experimental results implies control, which may have already been exercised or must accompany the transferability of the experimental results by either 'exporting the lab' or by 'adapting the target at hands'. As explained above, an experiment where control is weakly exercised is unintelligible whose results possess both weak internal and external validity.

Guala's reply to Siakantaris presents a more optimistic prospect:

Unlike most economic systems, auctions are relatively simple systems that work in relative isolation from other economic and non-economic systems mechanisms. This may very well be true, but these rare lucky examples provide us with clues regarding what can be done also in less neat situations. Like the strategy of establishing parallelism by exhaustion, eliminating one possible background factor after another, the above strategy of data-to-data comparison is certainly not the end of the story. Much more sophisticated methods can probably be devised to support parallelism inferences, methods applicable to a wider range of circumstances. (Guala 2002, p. 264)

In any case, the fact of the matter is that parallelism most likely applies to simple contexts that exist in relative isolation. In chapters 7 and 8 yet more optimistic prospects are presented. These stem from the observation, already alluded to above, that an economic experiment is a social context that provides relevant information about human behaviour and socioeconomic institutions. I will argue that the

information economic experiments provides extends beyond the contexts that generated. This is an intuition that is not captured by the notions of internal validity, external validity, or parallelism for that matter.

Again, the sharp contrast between the views on the internal and the external validity issues is perplexing. The almost unreflective belief in experiments as auxiliary tools for theory development coexists with disbelief about the generalisability of experimental results to nonlaboratory situations. Symptomatic of this is Starmer's words on the matter:

As experimentalists, for the most part we simply do not know whether the results of laboratory experiments apply more generally in everyday contexts of economic significance. The official theory-testing rhetoric of experimentalists attempts to insulate us from this fact. As experimentalists, then, let us admit this limit to our present knowledge and pursue a broader exploration of the extent to which laboratory findings generalize outside the laboratory. (Starmer 1999a, p. 25)

This anxiety is a manifestation of a theory-oriented view of science that takes 'reality', 'realisticness', and 'simulation of some natural occurring phenomenon' as the sole criteria for assessing the epistemic value of economic experiments. The social epistemology of experiment relieves the dependence of the epistemic appraisal of economic experiments on some existing referent in the real world. Insofar as experimental results are generated by the agency of experimental subjects, it provides knowledge about human behaviour and about the socioeconomic institutions *whether or not* the experimental situation finds some parallel in the real world.

6.6 THE SOCIAL DIMENSION OF EXPERIMENTAL ECONOMICS

The social epistemological argument of experiment alluded to the 'sociality' of experimental practice in economics. In this section the epistemic value of the collective nature of knowledge production in economics is analysed in more detail.

It was noted that economic experiments are amenable to the scrutiny of the larger scientific community, which allows extending the process of knowledge production and thereby the formulation of interesting problem-situations and the production of reliable problem-solutions. The analysis of the practice of experimental economics shows that this potential has been actualised. A characteristic trait of experimental economics is precisely the incremental way by which knowledge is generated. That is, in experimental economics knowledge is produced by series of experiments whereby economists build upon one another's results. This aspect of the experimental process of knowledge production attenuates many of the methodological problems identified above. The incremental way in which knowledge is produced is very clear in the surveys of experimental economics and it is explicitly remarked by Alvin E. Roth:

For the first time [in the 1950's], there began to be a wide variety of areas in which different groups of experimenters began to study the same issues from different points of view. This

¹³² This apparent limitation of experimentation is not specific to economics. Students of experimental physics have made the same point. They have argued that the successful application of an experimental result outside the laboratory requires the manipulation of the world so as to make it like the laboratory (Cartwright, 1999).

meant that there began to be *series* of experiments in which investigators with different hypotheses responded to one another's experiments, critically examining earlier conclusions. It is this process, in which experimental results suggest new experiments and in which different points of view suggest different experiments to different groups of experimenters, that allows us to begin to look back on experimental economics as a cumulative process. This kind of dialogue is one of the great sources of strength of the experimental method. (1995a, p. 21)

The intuition underlying this argument already appeared in Smith's response to the DQ thesis that undermined the epistemic value of a single experiment to corroborate or refute a theoretical hypothesis. It was argued then that if there is the suspicion that a given result is caused by some extraneous factor, this suspicion can be clarified by the design of further experiments to probe the source of the theory's failure or success. As Smith put it: "when a theory works well, they push imaginatively to find deliberately destructive experiments that will uncover its edges of validity... when a theory works poorly, they re-examine instructions for lack of clarity, increase experience level of subjects, try increase payoffs, and explore sources of 'error' in an attempt to find the limits of the falsifying conditions" (1994, p. 129). The point is that by conducting series of experiments, experimenters are able to increase the level of confidence on the experimental results.

This pattern of practice is pervasive in the surveys of experimental economics (e.g. Kagel and Roth 1995, Davis and Holt 1993, Roth 1988). In the early stages of a research program not much is known about the means and the results of experimental practice, and there is ample room for disagreement. As the sources of disagreement are identified and experiments conducted to elucidate them, the dissent gradually reduces. At first, follow-up experiments investigate whether or not the experimental results are to be attributed to an artefact of the experimental procedure. This normally calls for the re-examination of the precepts and procedures of experimental economics. But as the results become accepted, attention is redirected to the exploration of factors that may have an effect on the phenomenon under scrutiny. Only at a later stage, when the phenomenon is fully understood, experimenters try put forward and test tentative explanatory hypotheses. Earlier results are reinterpreted, areas of disagreement are narrowed down and what were apparently conflicting results may eventually be integrated into a larger unifying account.

[S]ubsequent experiments sometimes define more narrowly the conditions under which some initially observed phenomenon occurs and sometimes cause the results of an earlier experiment to be entirely reinterpreted and initial conclusions rejected... series of experiments can be constructed to allow us to draw more *reliable* conclusions, both about what we know and about what we know we don't know. (Roth 1995a, p. 23, emphasis in original)

Thus, and similarly to scientific experimentation in the natural sciences, the social dimension of knowledge production performs the epistemic role of controlling the constraining effect of extant scientific culture and experimenters' beliefs in knowledge production. As Roth notes,

I think the major pitfall to be aware of here is that, since decisions made in the design of the experiment cannot be regarded as random samples from the space of possible design choices, *there is room for an experimenter's prior beliefs about the likely outcome of the experiment to influence the outcome, through these design decisions...* The reason is that the designer of an experiment may have to make many particular decisions, and choose the level of many parameters that may affect the phenomena under study. If there is reason to believe that the resulting observations depend in important ways on some of these choices, variations may be incorporated into the experimental design, or into subsequent experiments, to see if this is the case. (1988, p. 1023, emphasis added)

The fact that knowledge production involves subjective beliefs does not necessarily jeopardize the experimental enterprise. Intersubjective disagreement can be an epistemic asset if it generates novel experiments intended to resolve the points of contention. This requires a critical community of scientists who evaluate and submit the results of science to further scrutiny. The critical quality of the scientific community, in turn, depends on the community's pool of prior commitments. As Roth put it, "experimentation is well served by skeptical readers and particularly by experimenters with different theoretical predispositions" (1988, p. 1024). To repeat, scientific community must be critical and its critical nature hinges on the distribution of beliefs within the community. To conclude, the incremental way whereby knowledge is generated constitutes a collective endeavour that attenuates the effect of consciously or unconsciously held beliefs on the experimental process of knowledge production.

6.7 THE EPISTEMIC PORTRAIT OF EXPERIMENTAL ECONOMICS

Experimental economics consists in the creation, control and measurement of microeconomic systems in the laboratory to generate knowledge useful for theoretical development or learning about some aspect of human behaviour or about socioeconomic institutions. The arguments presented by economists to justify the epistemic value of experimental microeconomics together with the social epistemology of experimental economics provide a very precise picture of economic experiments that is synthesised below.¹³³

1. THE LEVEL OF 'MATERIALITY' OF ECONOMIC EXPERIMENTS IS LOW.

The analysis of materiality involves assessing the material quality and the material agency of the experimental system. The *material quality* of an experimental system refers to the 'socioeconomic quality' and the 'human quality' of the experimental systems. The *material agency* of the experimental system refers to the 'socioeconomic agency' and the 'human agency', i.e. the contribution of the microeconomic institution and subjects' actions to the generation of experimental outcomes.

¹³³ Recall that this portrait concerns what has been referred to as the 'outcome-related' strand of experimental economics, as opposed to the 'behavioural' one (see n. 63, p. 97).

The level of ‘materiality’ of economic experiments is low because both the human quality and the human agency of economic experiments are low. The phenomenon created in the laboratory is a simplified, abstract and artificial version of an actual or potential real-world situation. Experimental subjects are students who are asked to perform abstract tasks in the unfamiliar environment of the laboratory. In addition, the laboratory is in itself a concrete socioeconomic context that departs in significant ways from non-laboratory situations. Indeed, it is this low level of socioeconomic and human quality that underlies the simplicity and the artificiality criticisms, especially so the latter.

The level of human agency is also low because the requirement of control severely constrains the agency of experimental subjects. In general, this is accomplished by inducing on the experimental subjects self-interest and income maximising motivations while directing the subjects’ attention to the performance of specific tasks. The control exercised over human motivations is necessary to guarantee the intelligibility of the experimental system, in particular, the relation between individual behaviour, institutional rules, and the microeconomy’s performance.

Even though the requirement of control constrains the ability of experimental participants in frustrating scientists’ prior expectations, which is the main source of epistemic value of the experimental process of knowledge production, it does not undermine its role altogether. But it does limit the application of the experimental method of economics. Experimental economics allows the study of individual behaviour in very constrained environments, in particular, the study of factors that affect self-interest maximising behaviour and/or the study of microeconomic institutions that presuppose this behaviour. Because the epistemic value of experimental results derives from the participation of human subjects in knowledge production, the analysis of this epistemic value of experiments requires determining the extent to which the experimental results are the outcome of their agency rather than the agency of the experimenter.

2. ECONOMIC EXPERIMENTS ARE STRINGENT EXPERIMENTAL SYSTEMS.

Economic experiments are stringent experimental systems that tend to produce intelligible results. The stringency of economics experiments is due to two main factors, the high level of control exercised over the microeconomic system, and the high dependence of the design of economic experiments on economic theory. These factors together render economic experiments fairly transparent systems.

The relation between experiment and theory contributed to the rapid stabilisation of experimental results and consequent accumulation of experimental knowledge. The fact that the first experiments were built upon known and well-established theoretical resources to provide specific answers to well-defined questions significantly contributed to the transparency of original experiments. Economic experiments soon became well-understood resources of practice that could be easily adopted in subsequent practice. The specificity of the problem-situations of economic experiments was accompanied by the rigidity of its procedures that established, at a very early stage, a critical set of resources that set the standards for the design, implementation and interpretation of

experiments. Experimental economics soon became a fertile domain for solving particular problem-situations that the researchers could solve, understand, and render comprehensible to the wider community.

From this it follows that the more a problem-situation can be framed in terms of a microeconomic problem-situation *a la* Smith, and therefore the more the precepts of experimental economics can be satisfied, the more transparent the experiment is and the higher the ability of the experimental system to generate conclusive results.

3. THE STYLISED FACTS OF EXPERIMENTAL ECONOMICS ARE THREE-DIMENSIONAL ENTITIES THAT INCLUDE THE ASSOCIATED EXPERIMENTAL PROCEDURES AND INSTRUMENTAL MODELS.

The fact that economic experiments are designed to provide answers to specific questions in very controlled environments renders the results of experiments highly dependent on the experimental procedure and the associated instrumental model. In contrast to physics, where experimental results easily gain autonomy from the experimental practices that generated them (e.g. the non-existence of quarks), in economics experimental results remain attached to the experimental procedure and the instrumental model that produced them (e.g. market equilibrium is a phenomenon that occurs in *the* double auction experiment).

4. EXPERIMENTAL ECONOMICS IS CONFRONTED WITH A TRADE-OFF BETWEEN ‘MATERIALITY’ AND STRINGENCY.

The low level of materiality is associated with a high level of stringency of economic experimental systems. Conversely, a high level of materiality is associated with a low level of stringency. The fact that experimenters can be effective in exercising control over experimental subjects means that experimenters can design highly transparent experimental systems that produce conclusive results. Failure in exercising control over the experimental subjects, in turn, carries a higher potential for generating ambiguous experimental systems. This tension may, however, be attenuated by the collective and the incremental nature of the experimental production process.

5. ECONOMIC EXPERIMENTS HAVE A HIGH POTENTIAL FOR TECHNOLOGICAL APPLICABILITY.

The low level of ‘materiality’ and the high level of rigidity of experimental systems confer a high level of technological applicability to economic experiments. The technological applicability of economic experiments amounts to the design of institutions that attempt to achieve the best alignment between individual actions and some desirable socioeconomic goal. In this sense, economic experiments can be perceived as devices that examine how best to control human motivations for intended purposes. Insofar as experimenters are successful in achieving this control in the laboratory, at least in principle they may also implement it in other contexts.

6. EXPERIMENTAL ECONOMICS IS A FIELD OF RESEARCH WITH A LIMITED RANGE OF APPLICATION.

The requirement of control (and the ensuing imposition of the precepts of nonsatiation, saliency, privacy and dominance) renders microeconomic experiments instrumental to applications that presuppose self-interested and income-maximising behaviour. This is attested by the experimental work developed thus far within the fields of general equilibrium, social choice, industrial organisation, game and voting theory.

7. EXPERIMENTAL RESULTS ARE SOCIALLY ROBUST RESULTS INsofar AS THEY ARE AMENABLE TO THE SCRUTINY OF A WIDE COMMUNITY OF SCIENTISTS.

The fact that a significant bulk of experimental practice is part of the written reports of experiments renders the experimental process of knowledge production available to the scrutiny of a wide community of scientists. This is a peculiar trait of economic experiments that attenuates some of its methodological problems. The collective aspect of the process of knowledge production can significantly contribute to the clarification of ambiguous and controversial results because experimenters can easily appraise and improve upon the work of their co-practitioners. The collective dimension of knowledge production can therefore minimise the problems entailed by the trade-off between 'materiality' and stringency. This means that the collective aspect of knowledge production may allow for a higher level of human agency because alternative interpretations can be discriminated by conducting more experiments in subsequent practice.

6.8 CONCLUSION

This chapter reviewed the explicit and implicit epistemological arguments that justify the use of the experimental method of economics. Not all of them satisfactorily account for the major methodological and epistemic features of economic experiments. The less unsatisfactory arguments are informed by philosophies of theory that do not have a direct bearing on the methodological problems of experimental economics. The arguments that drew on the actual practice of experimental economics were, however, more illuminating. These invoke, more or less explicitly, the notions of 'human agency' and 'sociality' which answer more effectively the methodological questions of economic experiments.

The analysis of experimental economics supports the view that economic experiments are stringent devices whose epistemic value is to be assessed in terms of human *quality* and *agency*. On this view, an economic experiment is epistemically valid if experimental subjects understand the experimental situation in the intended way and if they actively participate in the generation of the experimental results. Questions that revolve around the 'simplicity' or 'artificiality' of experimental systems do not capture this intuition. The epistemological study of economic

experiments must appraise the participation of the human and the social worlds in knowledge production in order to evaluate the extent to which experimental practice generated interesting problem-situations and robust problem-solutions.

The analysis carried out in this chapter shows that the social process of consultation and evaluation has a very significant role in experimental economics. It attenuates the methodological limitations of experimental economics whose resources are simultaneously characterised by high levels of stringency and low levels of materiality. Moreover, the collective process of social evaluation plays an important role in the revision of scientists' prior commitments.

Experimental economics is now an established field of research with its own institutions and standards of practice. Experimenters are no longer as pressured to justify the experimental method when arguing for their results. However, the prevailing confusion in the field suggests that the methodological and epistemic attributes of economic experiments still need to be better understood.

APPENDIX 3

THE SIMPLICITY AND THE ARTIFICIALITY CRITICISMS

THE SIMPLICITY CRITICISM OF THE FICTITIOUS SCEPTIC

‘You take a group of people – whatever people come most readily to hand – and trust them into an artificial and unfamiliar situation. You then describe the experiment, perhaps getting them to read through several pages of instructions, followed by some test questions and/or a few minutes “practice time”. Then you set them going, with no real opportunity to acclimatize to the situation or think things through, still less to take advice and plan some course(s) of action. Instead, they have to go straight into action...’

‘But the real world differs from your experimental environment in certain important ways. If I’m going to enter into negotiations about something that matters to me, I’ll think out my strategy and tactics carefully in advance. If I’m searching for a job, a house, a car or some other major item of expenditure, I’ll engage in a rather more varied and sophisticated search procedure than your experimenters allow. If I’m deciding how much insurance to buy, or thinking about whether to move from a job with a low variance income to a job with a bigger variance but higher expected income, I’ll lay out the alternatives carefully, think about the possible consequences, take advice if necessary, and weigh up the pros and cons until I reach a considered balanced judgment.’

‘And if we’re talking about not just about individuals or households, but about private companies, public corporations, political and industrial interest groups and government departments, I’ll go further. We’re talking about agencies run by highly experienced specialists who can, if need be, buy in additional consultancy from the relevant experts. These are not five-minute first-guessers: millions of pounds may be at stake, and your entire research budget is but a fraction of what they are prepared to spend to work out their optimal strategy.’

‘That’s what the real world is like, and what economists are – or should be – interested in is what actually happens out there. Our models may be abstractions but if they give us reasonable accounts of actual behaviour which stand up well enough to statistical econometric tests, are we really supposed to call those models into question simply on the basis of a few little experiments involving a motley collection

of novices spending an hour or two earning paltry sums of pocket-money? How can that begin to compare with the evidence of what experienced agents actually do in situations where their decisions really matter?' (Loomes 1991, p. 598)

THE ARTIFICIALITY CRITICISM OF THE FICTITIOUS SCEPTIC

'They know that you're running an experiment, and so they know you're looking for something. Despite your protestations, they may believe there is *a* right answer which will be rewarded. Or they may want to please you, or create a favourable impression of themselves. Or they may simply feel pressure to do *something*, and so they look for cues, and seize upon whatever the experimental design (consciously or unconsciously) suggests to them. At best, their behavior is a first approximation, vulnerable to all kinds of biases from extraneous sources and from the way the experiment is set up and the way the decision problems are framed.' (1991 p. 598 emphasis in original)

FACTORS RESPONSIBLE FOR THE ARTIFICIALITY OF ECONOMIC EXPERIMENTS

- (1) Experimental situations often project a gamelike atmosphere in which a "subject" may see himself as "matching wits" against the experimenter-designer of the game. Even with large payoffs, a subject may derive personal satisfaction from perceived "victories" which are not necessarily correlated with the performance indices used by the experimenter.
- (2) Experimental subjects are often cast in roles, such as "seller", "dealer", or even "monopolist", and the subject may act in accordance with his own (mis)perceptions of these roles rather than in accordance with other incentives which may have been incorporated into the situation.
- (3) "Real world" behavior has usually been learned over many trials or over many years. The relatively short time horizons of experiments cannot hope to capture more than the behavior of the most naïve and inexperienced actor who are found in the wider system.
- (4) Among the biological species, human beings are foremost in their capacity to control their own behavior through the implementation of abstract rules. Since these rules can be applied to a variety of different situations, human subjects usually carry many of them into the laboratory. Which of many possible alternative rules is to be applied in the experiment then depends on the background and experience of individual subjects and on their short-run interpretations of the nature of the situation. (Cross 1980, p. 404)

CHAPTER 7

HUMAN AGENCY (OR LACK THEREOF) IN ECONOMIC EXPERIMENTS

7.1. IS EXPERIMENTAL ECONOMICS EXPERIMENTAL?

In the first part of the present work experimentation was reconstructed as a scientific endeavour that aims at establishing a relation of coherence among the three components of the experimental system – the 'material procedure', the 'instrumental model' and the 'phenomenal model'. It was argued that the mutual support of the components of the experimental system provides scientists with reasons for believing the experimental results. The three-way coherence assures scientists that sense has been made of the world and their practice in it. It was also noted that the epistemic value of coherence is conferred to by the direct participation of the material world. By resisting to scientists' expectations, the material world promotes the revision of prior beliefs and thereby the generation of novel knowledge.

In chapters 5 and 6, the analysis of the experimental method of economics revealed that economic experiments possess a low level of 'human agency' – the 'material' of economics – that stems from the methodological necessity of creating transparent experimental systems.¹³⁴ The experimental procedures and the precepts of economics constrain the agency of experimental subjects by reducing human motivations to those presupposed by economic theory – self-interest and income-maximisation – and by restricting the available range of choice to a few options. The control of subjects' motivations and actions is intended to enhance the intelligibility of the experimental system, specifically, the relation between the 'environment', the 'institution' and the performance of the 'microeconomic system'. However, the requirement of transparency limits the potential of economic experiments for learning about human motivations and behaviour.

Experimental economics is, thus, informed by an unavoidable tension. On the one hand, the hallmark of scientific experimentation is the direct engagement of the subject matter in knowledge production. On the other hand, experimental economics seems to be condemned to an intrinsically reduced level of such participation due to the tight control over the motivations and actions of experimental subjects. This tension raises the obvious question: *What kind of knowledge does experimental economics provide given that its methodological precepts require tight control over human behaviour?* This chapter

¹³⁴ Recall that this characterisation applies to the 'outcome-oriented' kind of economic experiments.

provides the answer to this question by introducing a distinction between technological experiments and behavioural experiments.

The classification of an economic experiment as a technological experiment or a behavioural experiment is intended to emphasise the role human subjects play in the experimental process of knowledge production. *Technological experiments* are devices for institutional engineering. Their aim is to design microeconomic institutions that serve specific purposes. In these experiments human subjects are merely instrumental to that end. They allow testing the robustness of microeconomic institutions to human vulnerability and variability. The level of control over human motivations and actions is high and consequently the level of ‘human agency’ is low. Because the level of ‘human agency’ is low, the prospects for learning about human motivations and behaviour are significantly reduced. But the level of stringency is high insofar as technological experiments are fairly stable and transparent systems.

In contrast, *behavioural experiments* aim at learning about human behaviour in various socioeconomic contexts. Because the goal is to learn about human behaviour, control of human motivations and actions must be low. Or to put it in another way, the level of ‘human agency’ must be high so that human motivations may be manifested in the experimental setting and knowledge about them generated. The level of stringency is low because behavioural experiments are fairly plastic experimental systems amenable to various interpretations. This difficulty is attenuated by the collective nature of knowledge production that promotes follow-up experiments that eventually discriminate among the alternative explanations. The main argument of this chapter is that the high level of ‘human agency’ renders the behavioural experiments of economics experiments proper because they have good prospects of generating knowledge about individual and collective behaviour. The potential of technological experiments for generating novel knowledge about human behaviour is much lower, but they constitute a highly effective instrument for socioeconomic engineering.

Given that the previous chapters have addressed technological experiments rather extensively, after a brief summary of technological experiments in section 2, the remainder of chapter focuses exclusively on behavioural experiments. The goal is to demonstrate their potential for providing knowledge about human motivations and behaviour.

7.2 THE TECHNOLOGICAL EXPERIMENTS OF ECONOMICS

The market experiments reviewed in chapter 5 are technological experiments. Their goal is to evaluate the efficiency of various institutional arrangements under specific market conditions, which are defined in terms of type of commodity, number of traders, uncertainty, and so on. The rules by which subjects communicate and engage in exchange are the object of analysis. In these experiments subjects are more adequately perceived as a piece of the experimental apparatus rather than the object of study. The behaviour of the experimental subjects is manipulated in a known, regular, and reliable way for the study of market rules. Specifically, in market experiments subjects are instrumental to the appraisal of the informational, communicational and

exchange attributes of market institutions, the performances of which are evaluated on the basis of subjects’ actions under those rules.

The control exercised over subjects’ motivations is highly effective and stable across experiments. In chapter 5 it was shown that control is exercised by inducing self-interested and income-maximising behaviour via a reward structure that prizes economically successful decisions undertaken in a context of relative anonymity. Standardised experimental procedures, in addition, help implementing the precepts of experimental economics by shielding the microeconomic institution from the wider socioeconomic context.

The goal of market experiments is to study the performance of market institutions. Even though the characteristics of market institutions can only be identified and appraised by having subjects interacting under those rules, divergent performances implicate the rules of communication and exchange rather than human subjects. This, of course, is not to say that human behaviour is irrelevant. Both human behaviour and the institutional setting determine the performance of the microeconomic system. This is precisely where the significance of market experiments lies. The objective is to design rules of communication and exchange that achieve the best alignment between individuals’ preferences and actions and the system’s performance.

Nor does it follow that technological experiments are trivial. The control over human behaviour does not ensure perfect matches between individual actions and desirable aggregate outcomes. For one thing, even under an appropriate incentive structure individuals may fail to follow the course of action that best suits their interests. Generally, this occurs when the task is too complex so that it is difficult for subjects to identify the course of action that maximises income (e.g. the winners’ curse). For another, the best course of action from the individual viewpoint may produce undesirable collective outcomes (e.g. market collusion). The relevance of technological experiments therefore resides in this potential mismatch between individual preferences and individual behaviour, or between individual and collective goals. Their purpose is to design an institutional arrangement that guarantees compatibility between the interests of economic agents and socially desirable goals.

It follows from this that the ‘human agency’ of technological experiments manifests itself when experimenters face difficulties in establishing intended matches between individual behaviour and aggregate outcomes. However, the level of participation of experimental subjects in these experiments does not significantly promote the acquisition of knowledge about individual behaviour, at least not beyond the knowledge of how best to control it so as to avoid undesirable outcomes at the system’s level.

The low level of ‘human agency’ of technological experiments is associated with a high level of stringency. Technological experiments are rigid systems that operate in a stable, reliable and intelligible way under varied conditions. Because they are successful in inducing behaviour in laboratory contexts, they bear a potential to induce the same behaviour in similar non-laboratory contexts. This explains the technological success of these kinds of experiments, as mentioned in chapter 5. Hence the classification attributed to them.

7.3 THE BEHAVIOURAL EXPERIMENTS OF ECONOMICS

In contrast to technological experiments, behavioural experiments have a strong potential for generating knowledge about human behaviour (hence the label attributed to them). Given that the requirement of control is the obstacle to the generation of knowledge about human motivations and behaviour, it is not surprising that experiments that fall within this category have more difficulty in controlling them. And insofar as control is exercised by the precepts of experimental economics, neither is it surprising that control fails when the four precepts of experimental economics cannot be implemented for the purposes at hand. But as experimenters lose control over human behaviour, they enhance the chance of generating novel knowledge about it. To put it differently, by bringing about varied motivations, these experiments carry a higher potential for the growth of knowledge about human behaviour. Insofar as control in economics consists in inducing self-interested income maximising behaviour, the potential of behavioural experiments lies in the discovery of counteracting factors to self-interest and reward-maximisation. The high degree of plasticity of behavioural experiments, however, creates difficulties in interpreting the experimental results. This is a problem that can be solved by the collective aspect of knowledge production and by subsequently conducting novel experiments. How this is accomplished is exemplified here with the ultimatum game experiment.

The ultimatum game experiment is an exemplar of a behavioural experiment that presents the virtues of simplicity while maintaining a high degree of affinity to the technological market experiment. It therefore is a useful foil against which the market experiments may be compared. The aim is to show how the superior level of ‘human agency’ of behavioural experiments entails a higher potential for the generation of knowledge about human motivations and behaviour.

The first ultimatum experiment was conducted by Werner Güth, Rolf Schmittberger and Bernd Schwarz (1982). Subjects were randomly divided into two groups and placed in opposite sides of the same room.¹³⁵ One group was composed of subjects that would perform the role of player 1 and the other of subjects that would perform the role of player 2. Each player 1 would engage in a two-round game with a player 2 in which nobody would know with whom they were playing with. In the first round, player 1, who comes to be referred to as the proposer, would determine how to distribute a fixed amount of money. In the second round player 2, later known as the responder, would decide whether or not to accept the proposed distribution. If player 2 accepted, each received according to the proposed distribution, otherwise both received nothing.¹³⁶ Under these conditions a very asymmetric distribution should follow. The game theoretical prediction (assuming that each subject seeks to maximise his own payoffs) is that player 1 will receive the bulk of the fixed amount. Supposing that a rational player will always prefer the alternative that yields the higher payoff, player 1 should offer the

smallest positive payoff because he can expect that it is accepted by player 2, who accepts any positive offer rather than rejects it and earns nothing.

But the experimental results departed from the theoretical prediction. Not only did the proposers make more generous offers – only a few players demanded almost the totality of the dividends and the modal offer was actually the equal split – but the responders also refused positive payoffs by opting for no rewards. The authors concluded that these results constituted evidence for the fact that players care about the relative share in addition to caring about the absolute amount of their payoffs.

Güth, Schmittberger and Schwarz (from now on GSS) suggested that the failure of the proposer in exploiting the first-mover advantage could be explained by the experimental context which rendered unacceptable the exploitation of an advantageous position. The underlying presupposition was that the bargaining situation consists of a game between two opponents who would stand in an equal position were it not the arbitrary allocation of the different roles between them. GSS then ventured that an asymmetric relation would be more acceptable in a market context, such as, the consumer markets of industrialised countries where “buyers... might be used to have less strategic power”. In contrast, “in the abstract bargaining situation, where the bargaining parties have to divide a given amount of money, an asymmetric power relationship is probably less acceptable” (p. 369).¹³⁷

It is now possible to show the extent to which the ultimatum game is a behavioural experiment. Behavioural experiments allow acquiring knowledge about human behaviour while technological experiments attempt to control it for some specific purpose. This difference emerges very clearly from the analysis of the aims and outcomes of each kind of experiment. The market experiments aim at studying the performance of different market institutions for which end the experimental subjects have an instrumental value. Consequently, the significant data lies in the transaction prices and quantities on the basis of which the performance of the microeconomic system is evaluated. In contrast, the aim of the ultimatum game is to observe individual decisions in a particular strategic environment. Consequently, the relevant data consists in those decisions, that is, the proposed partitions and the rates of acceptance or rejection. Thus, while in the former experimental results bear on the institutional rules of the microeconomy, in the latter the results bear on the motivations and preferences underlying the observed behaviour – preference for equity, willingness to exploit an asymmetric situation and tolerance towards the exploitation of asymmetric power by another party.

The ‘human agency’ of the ultimatum game experiment is higher than that of the market experiment. As we have seen, in the double auction the level of ‘human agency’ is low because subjects are explicitly told to make the highest profit while the experimental context does not motivate subjects to do otherwise. Buyers are advised to bid below their reservation prices and sellers to make offers above theirs. And both buyers and sellers have no reason not to conform to this prescription. Moreover, given that the goods with higher reservation prices tend to be the first to

¹³⁵ The instructions are provided in Güth et al. (1982).

¹³⁶ Specifically, player 1 would write in a form the amount of money he demanded for himself which would be distributed by chance to the subjects of the other group. Player 2 would write down whether or not he accepted the proposal.

¹³⁷ This conjecture is based on the confrontation of these results to those of Lawrence Fouraker and Sidney Siegel’s bilateral monopoly experiment (1963) in which the seller is successful in exploiting his monopoly position.

be bought and the goods with the lower reservation prices tend to be the first to be sold, there is a trend for the market price and the market quantity to converge to the equilibrium levels (see Ch. 5). In the ultimatum game, proposers are asked to divide a fixed amount of money for which a wide range of options is available (from very unequal to equal partitions) and which may generate divergent reactions from the responders. In short, whereas in the market experiments there is one salient course of action common to all subjects, in the ultimatum game different courses of action are available depending on subjects' perceptions of the experimental situation.

The variety of available courses of action in the ultimatum game is strictly related with the violation of one of the methodological precepts of experimental economics – 'privacy'. 'Privacy' is not guaranteed in the ultimatum experiment due to the very nature of the game. Since the payoffs of the players are the very object of decision-making each player knows the payoffs of the player they are interacting with. Because intersubjective considerations cannot be avoided, the range of motivational factors at work may be high, and so may be the set of available courses of action. Moreover, intersubjective considerations may be further stimulated by the fact that the ultimatum game constitutes a more personal kind of interaction. It is a two-player game as opposed to the market context in which subjects address the whole group of traders. In sum, while in the market experiment subjects only know their own payoffs, which depend on their decisions and on those of the group as a whole, in the ultimatum experiment the payoffs are public information and depend on the details of the interactive process with another individual. Consequently, in the ultimatum game a greater array of motivations and preferences (e.g. self-centred or other-regarding preferences, tolerance towards unfair proposals,...) may be in operation which may give rise to different decisions.

The range of motivations may be further enhanced by a high degree of disparity in subjects' perceptions of the experimental context. On the one hand, the proposers' decisions depend on their views about what an adequate proposal is and on what they think the responders' expectations are. The responders, in turn, may have different views about the reasonability of the offer made. As a result, a variety of courses of action is possible and feasible. Consequently, the ultimatum game has a high potential for generating unpredicted behaviour and thereby knowledge about human behaviour.

The ultimatum game contrasts with the market experiments where subjects are not informed about their relative position because they do not know the others' reservation prices. Thus, they may believe that they stand in a fairly symmetric relation in the sense that they may enjoy equal opportunities to maximise their payoffs. Thus, the market context offers a rather impersonal and fair interactive context which does not present evident intersubjective considerations. Because the condition of privacy can be satisfied, eventual asymmetries may be unperceived by the subjects, and if perceived they may be evaluated as more acceptable in the market context.¹³⁸ Consequently, resistance towards inequality does not act as a strong countervailing motive to self-regarding and income-maximising behaviour.

¹³⁸ In line with GSS point, suppliers (consumers) may be more justified in taking advantage of a monopoly position insofar as they attempt to affect the price of the good for which they are the sole producer (buyer). In contrast, the ultimatum game offers a more abstract context where a position of privilege may be more easily perceived as undeserving.

7.4 EXPLORING HUMAN BEHAVIOUR

A wide range of experiments was carried out to investigate the effect of potentially relevant factors to GSS results. However, the results turned out to be robust to various modifications introduced to the experimental design. The original results of the ultimatum game were robust to: experience in multi-period games (Roth *et al.* 1991, Ochs and Roth 1989, Bolton and Zwick 1995, Slonim and Roth 1998), culture (Roth *et al.* 1991, Henrich *et al.* 2001, 2004)¹³⁹, and to the amount of the total payoff (Hoffman *et al.* 1996b; Cameron 1999; Slonim and Roth 1998).¹⁴⁰

The results of ultimatum game have become well-accepted stylised facts, which can be summarised as follows: (i) there are almost no offers below 10% or above 50% of the amount to be distributed; (ii) the modal and median offers are in the interval between 40%-50%; (iii) the means are around 30-40%; (iv) offers of 40-50% are rarely rejected; (v) offers below 20% are rejected about half the time; (vi) the rejection rate increases with the decrease of offers.

Because the results conflicted with the theoretical prediction, the ultimatum game caused perplexity. In particular, the experimenters were interested in understanding proposers' generous offers and responders' rejections of low but positive offers. Tentative interpretations were suggested and new experiments designed and implemented to test them. As mentioned in chapter 6, a characteristic feature of experimental economics is precisely this possibility of discriminating among conflicting interpretations by conducting more experiments.¹⁴¹

Two grand categories of explanations may be distinguished in this respect. The *fairness explanation* interpreted experimental results in terms of subjects' beliefs about the fairness of the proposal and the *strategic explanation* interpreted results in terms of proposers' expectations about the reaction of the responders.

The strategic explanation interpreted offers above the theoretical prediction as an indication of a correct anticipation of a high probability of refusals of extremely low offers because of the lower cost of rejection. Kenneth Binmore, Avner Shaked and John Sutton (1985) found evidence for this hypothesis in a two-period experiment where the cost of conflict of the first-period was increased by substantially reducing the amount to be distributed in the second period (by a discount factor of 25%). They concluded "the results indicate little support for the view that a substantial proportion of the population are "fairmen" as opposed to "gamesmen" (p. 1179). However, the theoretical prediction of this game represented a more balanced and fair distribution (an offer of 25%) than the extreme prediction of the original

¹³⁹ It should be noted that the study of Roth *et al.* involved four 'westernised' countries (USA, Japan, Slovenia and Israel). The cross-cultural study by Henrich *et al.* (2001, 2004) in fifteen small-scale societies of Latin America and Africa revealed sharper differences across cultures, with some societies closely conforming to the game-theoretic prediction while others exhibited extreme generous offers.

¹⁴⁰ For a systematic account of the main results of ultimatum game see Camerer (2003, Ch. 2) which include the analysis of the effect of other variables such as gender, age, race, academic major, beauty, etc.

¹⁴¹ Only a selected number of experiments will be presented here. The goal of this chapter is simply to illustrate the epistemic specificity of behavioural experiments. Surveys on ultimatum game experiments are, however, available in Thaler (1988a), Güth and Tietz (1990), Camerer and Thaler (1995), Roth (1995b), and Camerer (2003, Ch. 2).

ultimatum game.¹⁴² As Werner Güth and Richard Tietz noted, the experimental setup was not conclusive for it rendered “the game theoretical solution socially more acceptable” (in Roth 1995b, p. 259).

To show that, Güth and Tietz conducted a two-period experiment in which the discount factors (0,9 and 0,1) generated extreme predictions equivalent to those of the original game. Nonetheless, the actual offers were still reasonably balanced. More interestingly, they observed that a significant number of first-period rejections were followed by lower demands in the second period. This means that players 2 prefer to refuse a higher but inequitable reward in favour of a lower but more equitable one. Jack Ochs and Alvin Roth (1989) subsequently replicated these results. They observed that in fact the majority of rejections by player 2 are followed by counterproposals of smaller amounts of money (81% of rejections result in such disadvantageous responses).¹⁴³ These results taken together support GSS’s original conjecture – responders concern not only about their monetary pay-offs but they concern also about their relative share. Thus, this set of experiments support the view that subjects’ perceptions of fairness, which involve comparing their share of the available wealth to that of the opponent, do have a role in the observed behaviour. But a strategic reading of the behaviour of the proposer was still possible. It could still be argued that generous proposals were strategic in the sense that they entailed a higher chance of positive gains.

Further experiments were carried out to discriminate whether the results constituted evidence for strategic behaviour on the part of the subjects or instead an indication of a preference for fairness. The ‘crucial’ experiment was conducted by Robert Forsythe, Joel Horowitz, N. Savin and Martin Sefton (1994) that confronted the ultimatum with the dictator game.¹⁴⁴ In contrast to the ultimatum game, in the dictator game player 2 cannot reject the proposal of player 1, who determines the payoffs of both. Hence, the decision of player 1 can only be interpreted as being the outcome of his or her preferences.

As expected, the experimenters observed more generous offers in the ultimatum game than in the dictator game. The modal offer was the equal split in the former and the theoretical prediction in the latter, that is, player 1 kept the totality of the amount for him or herself. In the dictator game 36% of the players made zero offers whereas none of the players of the ultimatum game did so. Nonetheless, they continued to observe a significant number (around 20%) of generous offers (up to 50 %) in the dictator game, which was an indication that some subjects are genuinely motivated by considerations of fairness. They concluded that fairness is not the primary explanation of the motivations of player 1 in the ultimatum game even though some subjects may be so motivated to various degrees.

This brief survey of follow-up experiments to the ultimatum game shows how the interpretation of the experimental results of behavioural experiments promotes the investigation of preferences and motivations underlying human behaviour. It also shows how this investigation builds upon previous experiments until more decisive designs are implemented and more conclusive results obtained. The conclusiveness of the results, in turn, depends on the disparity of the views in confrontation. The more divergent the views, the sharper the designs that attempt to discriminate among them.

Another set of experiments was carried out to investigate how considerations of fairness interact with strategy in bargaining situations. To that effect, Vesna Prasnikar and Alvin Roth (1992) ran a best-shot experiment and an ultimatum experiment which entail extreme equilibrium predictions but had generated conflicting results, unequal distributions of income in the former and relatively equitable distributions in the latter.

The best-shot experiment is a ten-period public good experiment in which player 1 first states a quantity q_1 after which player 2, knowing q_1 , states the quantity q_2 he or she is willing to supply. The amount of the public good is given by the maximum of these two quantities (either q_1 or q_2), and the payoff of each player is determined by subtracting the cost of production of the units supplied from the redemption value of the public good. The incentive structure of the game dictates that player 1 chooses $q_1 = 0$ and player 2 chooses $q_2 = 4$, which yields a payoff of \$3,70 for the former and a payoff of \$0,42 for the latter. This experiment was performed under two information conditions. In the complete information condition subjects knew the payoffs of both players, and under the incomplete information condition subjects only knew their own reward structure. The ultimatum game is a ten-period experiment which followed the standard two-rounds game in which player 1 had to decide how to distribute \$10 and player 2 whether or not to accept the proposed distribution. In both experiments players would keep the same role and interact with different players in each period.

Even though the predictions entailed extreme distributions in both experiments, in the best-shot experiment the choices converged to the predicted equilibrium whereas in the ultimatum game they departed from it. Moreover, in the best-shot experiment convergence was quicker in the full-information condition than in the incomplete-information condition. But more important to their purposes was the observation that in both experiments players 1 played strategically. In the ultimatum game players 1 did better (i.e. their payoffs increased) by deviating from the equilibrium (i.e. by increasing their offers) while in the best-shot game they did better (i.e. their payoffs increased) by converging towards the equilibrium (i.e. by decreasing their offers). In the experimenters’ own words:

Once player 1 departs from equilibrium and offers a positive q_1 , player 2’s best response is always to provide $q_2 = 0$. This conforms with the observation... that player 1’s got little positive reinforcement for departures from equilibrium. This contrasts with the incentives to player 2’s in the ultimatum game, who have more incentive to accept an offer the farther it is from equilibrium. (1992, p. 880)¹⁴⁵

¹⁴² Because the amount of money to be distributed reduced from 100 pence in the first round to 25 pence in the second round, the prediction was that player 1 makes a demand in the range 74-76 pence, which is accepted by player 2, who cannot do better in the second stage.

¹⁴³ The authors calculated the rates for the previous experiments which were around 75% in the Binmore et al. (1985) and 88% in the Güth et al. (1982).

¹⁴⁴ The first dictator experiment was performed by Kahneman et al. (1986). However, this was not directly comparable with the ultimatum experiment because this was a choice experiment between two proposals.

Having shown that the behaviour of the subjects accords well with the incentive structure of the experiments, the authors concluded that subjects' decisions in the ultimatum game are primarily strategic. Fairness considerations matter because they have strategic value to the subjects which is learned by the losses subjects incur when they make low offers. Of course, this interpretation only succeeds in explaining the behaviour of player 1. No strategic interpretation explains the rejection of positive offers by players 2 insofar as subjects change partners between periods. This means that the rejection of unfair offers cannot be interpreted as the building of a reputation for future interactions.

However, it is not obvious that strategic considerations overrode concerns of fairness. The best-shot experiment significantly constrained the manifestation of fair behaviour. The reason is that the equal split was not very significant. On the one hand, the equal contribution to the public good was unattractive due to the low level of the payoffs to both subjects (the maximum amount they could both win was \$0, 42 with $q_1 = q_2 = 4$). On the other hand, unequal contributions had no significant impact on the disadvantaged party (with $q_1 = 4$ player 1 at worst wins \$0, 42, and with $q_1 = 0$ player 2 at least receives \$0, 42). That is, to any positive contribution of player 1 the choice of an equitable distribution by player 2 does not increase the pay-off of player 1, while it decreases the pay-off of player 2 (see n. 145). Under these circumstances, player 1, who had the first-mover advantage, chooses not to contribute and thereby to benefit himself from the unequal distribution. Since the equal split was not attractive, the main question the best-shot game raised concerned selecting the beneficiary of the unequal distribution. If player 1 did not take the opportunity then player 2 would.¹⁴⁶ This means that in the best-shot game inequality was substantially induced by design. In the ultimatum game, in contrast, inequality depended on the preferences of player 1 and on how he perceived the context of interaction. To repeat, the unequal distribution of income in the best-shot game is most adequately attributed to a feature of the experimental design rather than to subjects' preferences and motivations. Consequently, the claim that in this experiment strategic considerations overrode fairness considerations is not fully warranted. Güth and Tietz (1990) make a similar objection by pointing to a trade-off between equality and efficiency in the best-shot experiment, while it was non-existent in the ultimatum game:

Equal positive contributions in best shot games are obviously inefficient since one of the two contributions is completely useless. If sharing the burden of providing the public good is impossible, fairness considerations cannot be applied. Furthermore, the very obvious aspect of efficiency requires extreme payoff distributions. (p. 428)

¹⁴⁵ In the best-shot game the structure of game is such that the more of the public good player 1 provides, the less incentive player 2 has to contribute. The reason is that for high levels of provision by player 1 the additional supply of player 2 below q_1 increases the costs to player 2 without increasing the redemption value of the public good. That is, the provision of player 2 has no positive effect. On the other hand, a provision superior to q_1 by player 2 entails a lower payoff to himself than no provision. Thus, the more player 1 supplies, the greater is the incentive for player 2 to withdraw production.

¹⁴⁶ The fact that in the treatment with perfect information subjects reached the unequal distribution more quickly might be an indication that subjects more easily understood that the equal distribution was not an outcome worth trying.

In response, Prasnikar and Roth designed a ten-period market game experiment that consisted of one seller and nine buyers in which each buyer offered a price up to \$10 (the redemption value of the good). If the seller accepted the highest offer he or she would receive the corresponding price and the buyer of the highest bid would receive the difference between the \$10 and the bid price.¹⁴⁷ The other buyers would not earn anything. In this experiment, the equilibrium prediction was that buyers bid the maximum amount of either \$9,95 or \$10, which would be accepted by the seller. This is the case because buyers raise the chance of being the highest bidder by bidding \$9,95 or \$10. Below these amounts the buyer does not earn anything for there might always be a higher bid that might earn the difference between the redemption value of the good and the bid price.¹⁴⁸ The seller, in turn, should never reject the maximum bid. Therefore, the theoretical prediction of the market game experiment is an extreme equilibrium distribution.

In contrast to one-shot experiment, the results conformed to the equilibrium prediction and were also consistent with a positive correspondence between the equilibrium trajectory and the incentive structure of the experiment. That is, the buyers could not improve their earnings by deviating from the equilibrium prediction. The experimenters concluded:

Taken together, these results suggest that although equilibrium predictions may need to be modified to take into account nonmonetary aspects of players' preferences (e.g., in the ultimatum games), nevertheless, even when equilibrium yields very unequal payoffs, strategic considerations are not *displaced* by considerations of equity. On the contrary, the best-shot and market games show that whether equilibria will be observed depends on the off-the-equilibrium-path behavior, which responds to the off-equilibrium-path incentives. (1992, p. 886 emphasis in original)

But, again, the results are not as self-evident as Prasnikar and Roth seem to suggest. Despite the fact that the equal split would be an efficient outcome, the design of the market game experiment rendered it very unlikely. Similarly to the best-shot experiment, subjects' behaviour in the market game experiment was substantially constrained. The problem was that the equal distribution would require coordination among subjects not to tender bids above \$5. This possibility was excluded at the outset because the experiment prevented communication among the subjects. Or alternatively, it would require that all subjects preferred the equal split and that they all believed that others had such a preference as well, so that no one would bid more than \$5. But as other experiments revealed (particularly the dictator game), fairness considerations vary among subjects.

The equal distribution of income between the winning buyer and the seller would require coordination among buyers because none of them had sufficient bargaining power to impose his or her preferred partition. In contrast to the ultimatum game experiment, where the proposer could define the partition, in the market game the offer

¹⁴⁷ In case of a match a random procedure would select the winner among the highest bidders.

¹⁴⁸ The bids are multiples of \$0,05.

was determined by the interaction of nine players who competed among themselves. In the face of competition, buyers were inevitably forced to tender attractive offers to raise their chance of winning. Thus, in this experiment buyers' behaviour cannot adequately reflect subjects' views on what a fair proposal is. It is more plausibly affected by the competition among subjects and their desire to be the winning bidder. In addition, the market context may have changed the experimental setting by rendering inequality between buyers and sellers more acceptable (see n. 137 and 138).

Subjects' behaviour may be more easily explained by subjects' reactions to the actions of other competitors than by their views on fairness. Under this interpretation, the observed behaviour would be motivated by a varied set of motives not present in other experiments, such as resentment and retaliation towards other competitors, which would account for the escalation of bids. Not only does the market game experiment render the equal split unfeasible, the unequal payoffs observed might have been affected by a range of motives different from fairness considerations. In fact, this is recognised by the experimenters:

Consider a hypothetical buyer whose preference for equality is such that his very first choice outcome would be to have all buyers submit identical bids of \$5 (or \$1), and who bids accordingly in the first two rounds. When he sees how high the actual transaction price is, he becomes annoyed with the other buyers, and (with the same motivation that would have caused him to express his displeasure by rejecting too small an offer if he were a seller in the ultimatum game) he decides to become the high bidder in round 3, in order to deprive other buyers of the benefits of what he sees as their unreasonable behavior. The point in considering such a hypothetical buyer is to observe that in *this* game his nonmonetary preferences cause him to behave in a manner indistinguishable from an income maximizer, while in the ultimatum game his preferences lead away from the equilibrium predicted for income maximizers. (Prasnikar and Roth 1992, p. 885 emphasis in original)

Therefore, the assertion that in the market game experiment strategic motives crowd out fairness considerations is not warranted. The extreme outcomes rather than being the result of 'strategic' thinking, it is instead induced by design that renders the equal payoffs unattainable. Indeed, there is some evidence for a preference for more balanced income distributions. In the best-shot there are zero contributions by player 2 when player 1 does not contribute (40 percent of the outcomes in which $q_1 = 0$ resulted in player 2 providing $q_2 = 0$) and in the market experiment there are \$5 offers despite evidence that these cannot be winning bids.¹⁴⁹

In conclusion, neither the best-shot nor the market game experiment are adequate designs to assess whether and to what extent a preference for fairness is displaced by strategic considerations. In the former, fairness is not a salient motivation because inequality is to be attributed to the conditions of the game that renders the equal distribution an unattractive outcome. In the latter, equality is unattainable due to buyers' reduced bargaining power. In the ultimatum game the equal split is both feasible and significant. Furthermore, in this game fairness considerations can explain the behaviour

of both the proposer, who makes generous offers, and the responder, who rejects mean offers. To put it in another way, these experiments cannot provide conclusive answers regarding the interaction between 'strategic' considerations and 'fairness' considerations because the level of 'human agency' is not comparable. In the ultimatum game experiment 'human agency' is high because human behaviour depends on the subjects' motivations and actions. In the other experiments, 'human agency' is low because the observed behaviour is strongly determined by the design of the experiments.

The foregoing analysis reveals the importance of assessing the level of 'human agency' when confronting the results of behavioural experiments. The elimination of surprising behaviour by limiting the range of interesting options to experimental subjects is not a very compelling result. It only shows how a particular pattern of behaviour can be eradicated. To summarise, the appraisal of behavioural experiments requires examining the level of 'human agency' and their comparative analysis requires examining their relative levels of 'human agency'. In other words, the commensurability of behavioural experiments requires equivalent levels of 'human agency'.

Ultimatum game experiments show that experimenters are not fully aware of the importance of analysing the actual contribution of experimental subjects, or that of the experimental set-up, to experimental results. The fact that experimental economics involve the participation of human subjects does not warrant the conclusion that the observed behaviour is the result of subjects' preferences and motivations. This has to be carefully ascertained when interpreting the results of behavioural experiments.

Before moving to the next section, it should be noted that the ultimatum game experiment has stimulated extensive research devoted to the study of the factors that affect the prominence of fairness considerations. Briefly, and focusing on the proposer's behaviour, Elizabeth Hoffman, Kevin McCabe and Vernon Smith (1996a, 1999) have shown that 'social distance' between subjects' affects their generosity. That is, their experiments have shown that the amount of offers increases with the proximity between the subjects (and the experimenter), and vice versa.¹⁵⁰ Similarly, Iris Bohnet and Bruno Frey (1999) demonstrated that knowing the identity of the recipient raises the amount of the offers. Sally Blount (1995) found that responders (or 'deciders' as she calls them) more likely accept small or uneven offers when these are generated by a random device rather than when these are the result of the decision of an interested party. Armin Falk, Ernst Fehr and Urs Fischbacher (2003) found evidence to the same effect in a choice experiment (in which the equal distribution was not always available). This experiment demonstrated that the rejection rates of intentional uneven offers are higher than just as uneven but non-intentional offers. Finally, the experiment by Andrew Schotter, Avi Weiss and Inigo Zapater (1996) demonstrated that when the right to make the offer is earned, not only does the proposer make lower offers but these are also more easily accepted by the responder (cf. Hoffman *et al.* 1994).¹⁵¹

¹⁵⁰ They designed a dictator experiment in which social distance is manipulated by varying the degree of the dictator's anonymity. The presupposition is that social isolation weakens the norm of reciprocity and therefore generosity. Indeed, in the maximum isolation condition there were 64% of zero offers, whereas in the minimum isolation condition there were only 18% of zero offers. Bolton and Zwick (1995) did not find such an effect in the ultimatum game.

¹⁵¹ See Camerer (2003, p.74-83) for a more complete account of this strand of experiments.

¹⁴⁹ The equal-payoff bid was the modal bid in four out of twenty market rounds.

7.5 EXPLAINING HUMAN BEHAVIOUR

Having accepted that self-interested income-maximisation cannot account for the experimental results, the economists set out to develop theories that incorporated other motivations and preferences. One such endeavour explained the experimental results in terms of ‘inequality aversion’. Ernst Fehr and Klaus M. Schmidt (1999) built a model in which fairness is understood as ‘self-centred inequity aversion’, that is, aversion towards a position of disadvantage that is unfavourable to oneself rather than aversion towards inequality in absolute terms. This model explains income distribution as the result of the interplay between the distribution of preferences within the population and the context of interaction. Specifically, the model predicts that the economic environment determines the preference type that is decisive for the prevailing behaviour in equilibrium. For instance, it shows that in environments with strong competition fairness considerations are ineffective even if there is only one player who does not prefer the equal split. In the case of proposers competition, each player has an incentive to make substantial proposals in order to turn part of the inequality to his or her advantage and thereby increase the chance of a positive payoff. In the case of responders competition, responders try to do so by accepting ever lower offers. Thus, under highly competitive environments fairness considerations are irrelevant. The aggregate outcome does not depend on them nor does it depend on their relative distribution in the population. Thus, as argued before, the observation of self-interested behaviour in some environments does not mean that fairness considerations were displaced by strategic ones, but the irrelevance of acting fairly therein. As Fehr and Schmit put it:

[T]he crucial observation in this game is that *no single player can enforce an equitable outcome*. Given that there will be inequality anyway, each proposer has a strong incentive to outbid his competitors in order to turn part of the inequality to his advantage and to increase his own monetary payoff. A similar force is at work in the market game with responder competition. As long as there is at least one responder who accepts everything, no other responder can prevent an inequitable outcome. Therefore, even very inequity-averse responders try to turn part of the unavoidable inequality into inequality to their advantage by accepting low offers. It is, thus, the impossibility of preventing inequitable outcomes by individual players that renders inequity aversion unimportant in equilibrium. (p. 834, emphasis in original)

Conversely, there are circumstances in which the existence of a few inequality-averse players creates incentives for a majority of purely selfish types to behave non-selfishly. But this requires that the inequality-averse subjects be willing to sacrifice some part of their material payoff to enforce equitable outcomes.

On a similar line, Gary E. Bolton and Axel Ockenfels (2000) explain a wide range of experimental behaviour in terms of a preference for equity. However, in their model, individual utility depends on both the absolute amount and the relative share of payoffs to the average payoff within the population. This means that behaviour is now interpreted in terms of aversion towards income inequality in a broad sense rather than unfair behaviour towards oneself. But neither the Bolton-Ockenfels model nor the

Fehr-Schmidt model can distinguish a rejection of uneven offers due to distaste for inequality from the desire to punish a player who made an unfair offer. That is, these models do not distinguish individuals’ reaction to unequal outcomes from individuals’ reaction to the processes that led to those outcomes.

A second strand of models does take the processes generating aggregate outcomes into account. Matthew Rabin (1993) interprets experimental results in terms of ‘reciprocal altruism’. The main idea is that “people do not seek uniformly to help other people. It is rather that they do so according to how generous these other people are. Indeed, the same people who are altruistic to other altruistic people are also motivated to hurt those who hurt them”. Consequently, “[i]f somebody is being nice to you, fairness dictates that you be nice to him. If somebody is being mean to you, fairness allows – and vindictiveness dictates – that you be mean to him” (p. 1281, emphasis omitted). Behaviour is then explained in terms of player’s beliefs about the intentions of other players. An intention is perceived as unfair if the action of the other player gives less than the equitable material payoff. For instance, in the ultimatum game, a responder rejects an unfair offer because he or she believes that the proposer has the intention to be unfair to him. This then triggers the desire to retaliate. The wish to punish the proposer is what justifies the willingness to support the cost of retaliating. It also explains why responders more easily accept uneven but non-intended offers that were generated by random procedures (cf. Blount, 1995). A similar model by Levine (1998) highlights the importance of judgments about each other’s altruism. Depending on the context, a relatively altruistic individual may behave meanly to a player believed to be spiteful and vice-versa. Thus, in some situations “regardless of how altruism is distributed in the population, there exist equilibria in which the coefficient of altruism does not matter, and that consequently these equilibria are the same as the equilibria of the selfish model” (p. 596).

A less formal suggestion by Colin Camerer and Richard H. Thaler (1995) interprets experimental behaviour in terms of general rules of conduct, ‘manners’, that apply to any situation. Noting the contradictory results of the dictator and the ultimatum game – in the former some dictators offer equal shares and in the latter responders inflict a high cost to the proposers by rejecting positive offers – the economists suggest that fair behaviour does not hinge on intersubjective comparisons. Because a “player does not care about the others’ welfare per se, but desires some kind of equity in the context of this particular interaction”, the outcomes of these games “have more to do with manners than altruism” (p. 216). The general principle is that: “[e]tiquette may require you to share a windfall with a friend, but it certainly does not require you to give up some of your hard-earned year-end bonus to a stranger” (ibid). Because in these games subjects feel that they are on a relative equal standing, the advantageous position of the proposer should not be exploited by him. If it is, rudeness is followed by punishment. In the market game with responder competition players do not have the power to punish the rude behaviour of the proposer. They can only punish the unfair behaviour of the players with whom they compete. In fact, unfairness may not even be attributed to the proposer because “the Proposer is merely accepting the best deal available”. It is more adequately attributed to the behaviour of the other responders “who are misbehaving

by accepting such low offers” and who can be punished by the subsequent acceptance of lower offers (p. 217).

A fourth account offers an evolutionary interpretation of the ultimatum game. On this view subjects have no concern about each other’s payoffs. The experimental results reveal adaptive behaviour. The central idea is that “pure strategies which have been played and have met with success tend over time to be played with greater frequency than those which have met with less success” (Roth and Erev 1995, p. 172). Thus, proposers learn to make fair offers because they learn with a cost that responders reject low offers. Responders reject low offers because they are not so costly. Moreover, as players 1 increase their offers, responders have a better incentive to accept them. The upshot is that the incentives of the interactive context, the low cost of rejection to the responders and the high cost of rejection to the proposers are effective in inducing fairness. Other evolutionary models explain reciprocity in terms of a predisposition that some individuals have “to punish those who violate group-beneficial norms, even when this reduces their fitness relative to other group members” (Bowles and Gintis 2004, p. 17). This is a natural predisposition that cannot be cancelled out in unnatural one-shot games with strangers.¹⁵² These are deficient explanations of the experimental data, however. To point just a few problems related with evolutionary theories, neither the psychological nor the cultural theories can explain the differences encountered between the one-shot and the repeated ultimatum games, or the differences between young children and adults, or the differences between societies. If individuals possess a natural disposition to reinforce beneficial group behaviour, than this disposition should be manifested in one-shot games, in different age groups and in different societies.

To summarise, the accounts reviewed here are overlapping and complementary explanations of the results of the ultimatum game experiments. They identify a wide range of explanatory factors that account or may account for them. These include social preferences, genetic or cultural propensities and strategic thinking. But the point I would like to make is that they all carry novel insights about human behaviour that are now starting to be incorporated in economic theory. These insights may be synthesised as follows:

Preferences vary among the population. There is a small proportion of pure altruists and of pure egoists, while the majority of the population concern about fairness. That is, in general, individuals care about being treated fairly, and know that others do so as well.

When placed in an interactive context, individuals form the expectation that unfair behaviour may be punished and that fair behaviour may be rewarded. Consequently, fairness considerations may be relevant irrespective of the individuals’ own preferences. The effect of fairness considerations depends instead on the possibility of some individuals expressing them by punishing or rewarding.

The fact that individuals share dividends in contexts where social distance is high and no opportunity for punishing or rewarding exists suggests, however, that there is a general sense of

fairness that guides human interactions. Perhaps, the same sense that motivates treating strangers fairly even with a reasonable amount of self-sacrifice.

7.6 CONCLUSION

This chapter contrasted two broad categories of economic experiments according to the role subjects play in the experimental process of knowledge production. *Technological experiments* use experimental subjects to study microeconomic institutions. The ultimate goal is to learn how to design institutional rules that best control human behaviour for the attainment of pre-specified socioeconomic goals. Insofar as the possibility of controlling human behaviour via the implementation of institutional rules can be extended to non-laboratory contexts, these experiments have a high potential for institutional engineering. *Behavioural experiments* use subjects to study human behaviour. Insofar as control in economics is achieved by inducing self-regarding and income-maximising motivations, behavioural experiments investigate the factors that may counteract them.

The classification of experiments highlights two epistemic goals of economic experiments: learning how to design microeconomic institutions and learning about individual motivations, preferences and behaviour. Of course, these are the poles of a continuum that characterise most of the experimental work of economics. The point I would like to make nonetheless remains. The more experimenters succeed in controlling subjects’ behaviour the more they avoid the resistance of ‘human agency’ and thereby the less will they learn about human behaviour. Conversely, the more an experiment possesses the characteristics of behavioural experiments, the higher is its potential to generate knowledge about human preferences and behaviour.

The distinctive role of economic experiments has methodological and epistemological implications. Insofar as behavioural experiments aim at drawing inferences about human preferences and motivations, the design and the evaluation of these experiments require the analysis of the level of ‘human agency’. The contingency of human motivations on the experimental context, and the constraining effect of experimental procedures on subjects’ actions, require examining the extent to which observed behaviour is the result of subject agency. This is a central topic to experimental economics which has not yet been given due attention. Technological experiments demand instead the appraisal of the extent to which human behaviour can be equally controlled in non-laboratory environments.

¹⁵² Other contributions on the same line include Güth and Yaari (1992) and Samuelson (2001).

CHAPTER 8

'HUMAN AGENCY' AND INFERENCE

8.1 INTRODUCTION

The previous chapter presented a taxonomy of experiments that distinguished technological experiments from behavioural experiments. It highlighted the differentiated epistemic role experimental subjects play in each category of experiment. This chapter derives the implications of experimenting with human subjects to the inferences that can be derived from laboratory settings. It identifies an intermediate kind of inference that has not yet been explored by the students of experimental economics. Finally, it distinguishes the content of the intermediate level of inference in technological experiments and in behavioural experiments.

Sections 2 and 3 compare the proposed taxonomy of experiments with other classifications. Section 4 presents how the participation of human subjects in laboratory experiments has been dealt with in methodological discussions. Section 5 puts forward the intermediate kind of inference. Section 6 brings the contributions of modelling to the discussion and section 7 illustrates how the intermediate mode of inference provides understanding about real world situations. Section 8 concludes.

8.2 THE METHODOLOGICAL CLASSIFICATIONS

The most common criteria for classifying economic experiments are their methodological functions. Even though different taxonomies have been proposed, the classification of experiments revolves around three main purposes: 1) theoretical testing and development, 2) production of 'stylised facts', and 3) recommendations for policy making.

Borrowing Kaplan's taxonomy, Smith distinguished between *nomothetic experiments*, *heuristic experiments*, and *boundary experiments* (1982, p. 940-2). Nomothetic experiments establish 'laws' of behaviour which may be inspired by theory (*nomo-theoretical experiments*) or by observed empirical regularities in the field (*nomo-empirical experiments*). Heuristic experiments explore new topics of inquiry. The main difference between these two is that nomothetic experiments employ "replication and rigorous control to reduce error" and therefore "provide the most compelling and objective means by which each of us, as scientists, comes to see what others see, and by which, together, we become sure of what it is that we think that we know" (op. cit., p. 940).

Heuristic experiments “are less likely to follow a rigorous design pattern than nomothetic experiments because (a) the objectives may not be as sharply defined by theory or by a hypothesized pattern which is thought to characterize previous experiments, and (b) the procedural mechanics of the experiment may be new and untested” (op. cit., p. 941-2). These experiments are, however, “an important source of new discoveries” because “[it] is through exploratory probes of new phenomena that attention may be redirected, old belief systems may be reexamined, and new scientific questions may be asked” (op. cit., p. 942). Finally, boundary experiments test the boundary conditions of behavioural laws. Therefore, they “have an important function in establishing the limits of generality of a theory, and setting the stage for important new extensions in theory” (op. cit., p. 942).

Along the same line of reasoning, but focusing on the experiments’ intended audience, Alvin E. Roth (1995a, p. 21-3) classifies three different kinds of experiments: *speaking to theorists*, *searching for facts* and *whispering in the ears of princes*. Experiments that speak to theorists produce results that intend “to feedback into theoretical literature”. Experiments that search for facts attempt to study “the effects of variables about which existing theory may have little to say”. Thus, whereas the former establishes dialogue with theoretical economists, the latter “are part of the dialogue that experimenters carry on with one another” (p. 22). These may be followed by experiments that “search for meaning”, that is, experiments that test proposed explanations of the observed facts. Once understood, the discovered facts may then feedback into theory. Finally, experiments that “whisper in the ears of princes” engage experimenters and policymakers. According to Roth, “[t]heir characteristic feature is that the experimental environment is designed to closely resemble, in certain respects, the naturally occurring environment that is the focus of interest for the policy purposes at hand”. This corresponds to “the traditional responsibilities of economists, to formulate advice on questions of policy” (ibid).

Similarly, Douglas D. Davis and Charles A. Holt (1993) discriminates the theoretically oriented *tests of behavioral hypotheses*, the empirically guided *search for empirical regularities*, and the *theory stress tests* that try to bridge “the gap between laboratory and naturally occurring markets” (p. 19-20).

Daniel Friedman and Shyam Sunder identify as purposes of experiments “to discover empirical regularities in areas for which existing theory has little to say”, and “testing theories” or “seeing which theories best organize the data”. But they also include “to measure individual (innate or “home-grown”) characteristics in the population” and “test-bed experiments” as a new instrument of institutional engineering (1994, p. 7-9).

The epistemological taxonomy presented in chapter 7 is not incompatible with these methodological classifications. In fact, the epistemological taxonomy may be considered a higher-order scheme that subsumes all these methodological categories. The classification of economic experiments as test devices of theoretical hypotheses, as instruments for the collection of facts, or as tools for policy making, can be thought of as sub-classes of both technological and behavioural experiments. The difference between them is that while technological experiments produce knowledge about microeconomic institutions, behavioural experiments produce knowledge about individual motivations, preferences and behaviour. For instance (as it was shown at

length in Ch. 5) the technological market experiment, first conceived as a test device (of competitive price theory), evolved to become a benchmark experiment for the search of new facts (concerning microeconomic institutions) until it could be reliably used for the design of market institutions to be implemented in the field (e.g. FCT auctions). The behavioural ultimatum experiment (see Ch. 7), in turn, was initially motivated by ‘the search for facts’, but then evolved into tests of competing hypotheses. Even though behavioural experiments are not suitable to perform the role of ‘test-beds’, their results are certainly relevant for policy-making. They can help identifying the factors that counteract self-interested and income-maximising behaviour and therefore the factors affecting the efficacy of economic policies that presuppose them.¹⁵³

The methodological classification of experiments presented in seminal articles and handbooks aims at organising the practice of experimental economics according to the purposes economic experiments may serve. More recently, a new proposal emerged with other ends in mind. Robert Sugden’s distinction between *experiments as exhibits* and *experiments as tests of theories* (2005) is presented as part and parcel of a methodological argument. Sugden’s taxonomy, based on experimental practice in the field of individual decision-making, is part of a general argument that claims that the principles of experimental practice depends on whether experiments are serving as exhibits or as test of theories (p. 191).

An exhibit “is an experimental design which reliably induces some specific regularity (or ‘effect’, or ‘phenomenon’) in human behaviour”. The interest in an exhibit lies on the fact that “the regularity it captures contravenes – or at least, is not adequately explained by – some received theory of decision-making behaviour” (p. 191-2). A test of a deviation theory is, in turn, a test of a theory that predicts deviations from a received theory of decision-making behaviour.

It is not difficult to see that Sugden’s classification may be thought of as a sub-classification of behavioural experiments in the domain of individual decision-making. According to Sugden, the characteristic trait of exhibits is the generation of anomalies to received theories of decision-making that predict rational behaviour. As we have seen, the generation of anomalies means that experimental subjects behaved differently from the behaviour induced by the experiments’ incentive structure. Or to put it differently, the production of anomalies occurs in experiments with a high level of human agency. The high level of human agency derives from the presence of motivational and/or cognitive factors different from those postulated by received theory.

The fact that the behaviour of experimental subjects in behavioural experiments is informed by multiple factors is explicit in Sugden’s account of ‘exhibits’. ‘Exhibits’ purportedly express the operation of various causal mechanisms that differ from those postulated by received theory. Consequently, “the strongest and most robust deviations from standard theory will be found in those decision contexts in which these independent causal mechanisms happen to coincide, each inducing deviations in the

¹⁵³ There is already an extensive literature that has had recourse to the results of the behavioural kind of experiments to demonstrate the inefficacy, or rather, the damaging effect of policies that conceive human behaviour to be solely guided by self-interest and income-maximising motivations (e.g. Bowles 1998; Gintis 2000; Gintis et al. 2005).

same direction" (p. 294). By the same line of reasoning, 'deviation theories' are theories that have as the explanatory factor a 'causal mechanism', a preference pattern, or a decision-making process, that differs from those postulated by standard theory. What follows from this is that the 'tests of deviation theories' must allow for human agency, too. They must allow experimental subjects to behave differently from the behaviour induced by the experiments' incentive structure. In short, the tests of deviation theories are also 'behavioural experiments' because they must have a high level of human agency.

Sugden uses this classification to make three methodological points. First, Sugden argues that deviation theories can be subjected to 'severe testing' despite the fact that these are theories constructed to accommodate experimental evidence. The argument is that a "deviation theory typically predicts many systematic deviations from the received theory in decision problems which, at the time the theory is proposed, have not been investigated". Consequently, "the theory exposes itself to new tests within a well-defined domain of controlled experiments" (p. 297). Second, Sugden claims that 'the search for confirmation' is allowed in exhibits. It is a legitimate practice to investigate whether a particular violation of a received theory occurs where it is expected to occur. "[I]f that is the only claim that is made on behalf of an exhibit, the question of how it was discovered has no bearing on its significance" (p. 300). Third, Sugden points out that the 'tests of deviation theories' are significant regardless of the strength of the effect produced. These experiments necessarily yield weak effects because these are produced by isolated causal mechanisms. In contrast, 'exhibits' produce stronger effects because they produce multi-causal phenomena. In sum, the understanding of the results produced by exhibits requires a manifold of 'non-decisive' experiments each of which test the effect of an isolated causal mechanism.

Underlying Sugden's proposal is the concern that 'exhibits' and tests of 'deviation theories' may be vulnerable to charges of confirmation bias given that these experiments are constructed to generate the expected results. Sugden's response is that the process of discovery has no bearing on the significance of this kind of experiments. A deviation to standard theory is significant in its own right. The social epistemology recommends some caution, however. Just as technological experiments render 'received' behaviour fairly expected, behavioural experiments render 'deviant' behaviour unsurprising. The deviations generated by behavioural experiments are created in a context in which different motivations and/or different cognitive processes are more salient than self-regarding and income maximisation. And this salience must be born in mind when analysing these experiments. Granted that any violation of a theory is important in its own right, its significance depends on the effect of experimenters' 'confirmation biases' in the production of the experimental results. The appraisal of this significance requires the scrutiny of the level of human agency, that is, the actual participation of experimental subjects in the generation of the experimental results or the degree of determinacy of the experimental set-up. As we have seen in the previous chapter, the differentiated level of human agency rendered the positive offers of the dictator experiment more significant than those of the ultimatum experiment, and these more significant than the extreme offers of the market game experiment. However, it should be clear by now that the 'confirmation bias' is not higher in behavioural experiments than in technological experiments. Even though

technological experiments are not designed with the goal of producing a particular behaviour pattern, be it an 'anomaly', a 'fact', or an 'empirical regularity', they do constrain behaviour thereby producing behaviour patterns of their own. At any rate, the level of 'human agency' of behavioural experiments is higher than that of technological experiments.

8.3 THE EPISTEMOLOGICAL CLASSIFICATION

The epistemological classification that distinguishes technological experiments from behavioural experiments highlights the relation between the participation of human subjects in economic experiments and the content of experimental results. The high level of human agency of behavioural experiments explains their potential for generating knowledge about individual motivation, preferences and behaviour. The low level of human agency of technological experiments explains their success for the design of microeconomic institutions.

Mary Morgan's (2002, 2003, 2005) taxonomy, on a very similar line, highlights the mode and the degree of participation of the subject matter in various kinds of experiments. However, this taxonomy does not take into account the specificity of economic experiments and, consequently, it does not discriminate among the various kinds of economic experiments. In fact, it has been constructed by reference to experimental practice in the natural sciences (see section 4.2). Nonetheless, Morgan's account also takes the participation of the subject matter, the material world, in knowledge production as her key criterion.

As we have seen, Morgan distinguishes *ideal laboratory experiments*, *mathematical model experiments*, *virtually experiments* and *virtual experiments* using the material quality and the material agency of the experimental systems as criteria.¹⁵⁴ The ideal laboratory experiment is the scientific device with the highest level of materiality whereas the mathematical model experiment is at the lower end of a continuum of vicarious experiments, which may be classified either as virtually experiments or virtual experiments.

Morgan's taxonomy is very illuminating in revealing how the 'material' component may be embedded in different kinds of experiments. First, Morgan identifies the relation between the degree of participation of the object of inquiry ('material world') in knowledge production and scientists' capacity of controlling it. That is, the active participation of the subject of inquiry in the 'ideal laboratory experiment' is associated with a relatively lower degree of control. She explicitly states that "however ingenious the scientist, *the material world can only be controlled and manipulated to an extent*" (2003, p. 119, emphasis added). In contrast, the material world is absent in the 'mathematical model experiments' and therefore control is the most effective in them.

¹⁵⁴ Morgan uses the expression 'materiality' to refer to the ontological similarity between the subject matter of laboratory experiments and its referent, in both the natural and the human worlds. Thus, in this section the 'material' terminology applies also to experimental economics.

Second, Morgan suggests that the superior epistemic power of experiments as compared to that of models is conveyed by the fact that the former generate unknown results whereas the solutions of the latter comprise information that the modellers have intentionally built into them. The claim is that experimental results cannot possibly be known prior to practice, at the very least not fully, because they involve information that is not yet known. In the case of mathematical models, it is only the computational limitations of scientists that prevent them from knowing the outcomes. In her own words:

That material world experimental demonstration is necessary precisely because the resources for the result *we expect to find* are not necessarily present in the experimental setup: we might have the wrong account or theory about what will happen or our knowledge of the world might be seriously incomplete. In the mathematical model world, of course, we know the resources for the result that *we do find* because we built those resources into the model that constitutes the experimental setup. (2003, p. 220, emphasis in original)

Thus, the central thesis of the present work is also explicit in Morgan's account of experimentation: the direct participation of the subject matter in knowledge production (the 'material agency' in the natural sciences and the 'human agency' in the human sciences) conveys higher epistemic value to the results of experiment because it promotes the generation of novel knowledge.

Third, and more recently, Morgan came also to recognise that the epistemic value of experiments depend on the extent to which the experimental set-up constrain ('tame') the participation of the 'material world' (in the natural sciences) or the actions of the experimental participants (in the human sciences). In economic experiments

... much is determined and constrained by the design set-up and rules of experiments. This is clearly as it should be – there must be experimental controls... Laboratory experiments, by their design, investigate tamed economic behaviour..., that is by the experimental design and experimental rules that constrain behaviour. However, if the behaviour of those taking part in the laboratory experiment is entirely constrained, then the results will be determined absolutely by the experimental design and rules. Particularly when models are used within an experiment to provide measures of control, they may tighten the theoretical constraint beyond usefulness. They raise the danger of over-taming the participants in the particular way so that participants are no longer domesticated, but agents whose behaviour is directed by models of the world, models dictated by the economist. *An important question for any experimental design therefore is: Where is the potential for independent action in the experiment?* (2005, p. 325 emphasis added)

The social epistemology of experiment furnishes a normative framework for the appraisal of the 'potential for independent action' or 'human agency' in the generation of the experimental results. This is the task of the 'materiality' test (see section 4.2). In addition, the epistemological analysis of economics experiments must take into account whether the experiment is a technological or a behavioural experiment.

8.4 MATERIALITY, HUMAN AGENCY AND INFERENCE

The agency of the 'material world' and the agency of the 'human world' are the main factors that account for the epistemic superiority of laboratory experiments in the natural and in the human sciences, respectively. They are sources of knowledge that are absent in other processes of knowledge production. However, the possibility of extrapolating the information generated by experimental means to non-laboratorial contexts requires closer scrutiny. In particular, it requires the analysis of the relation between the experimental system and its referent in the real world.

The analysis of inferences from economic experiments to non-laboratory contexts has thus far been based on models. Models and experiments have similar functions in science. Indeed, models can be conceived of as experiments and experiments can be thought of as models (Mäki, 2005). Models can be understood as experiments whenever they are designed and manipulated in an experimental-like fashion (cf. section 4.2 above). Experiments can be thought of as models when they are designed to represent some real system in the wild or a theoretical one. As matter of fact, the hybrid nature of experiments, as depicted by Morgan, shows that most experiments possess both experimental and model attributes. The latter are more significantly present the more difficult is to experiment on a material representation of the object of interest so that quasi-material or non-material models are used in their place.

This analogy is particularly relevant when analysing inferences from the model or the experimental worlds to concrete situations in the real world. The simplified and representational characteristics of models and experiments are crucial to produce relevant information about the represented parts in the real world. On the one hand, experiments *qua* models are by necessity simpler model systems that allow the study of the more complex and intractable targets of the real world. On the other hand, if they are to generate information about the real world, experiments *qua* models 'resemble to' and are 'representative of' what it is that they represent (Mäki, 2001). That is, in order to adopt the role of representative of a target system, an experiment must represent it by resembling it in certain respects and to certain degrees depending on the purposes of the experiment.

The inferential power of experiments is higher than that of models because of the deeper ontological resemblance between experimental systems and real world systems. The fact that experimental systems are made of the same materials (natural or human) of the target they represent confers inferential superiority to experiments. Of course, this ontological similarity is variable. The inferential power declines as the media of experimental models deviate from the target system. For instance, the use of model organisms in biology allows the making of immediate inferences to other organisms of the same population and to extend these to other organisms with which they establish a relation of similarity. As Morgan puts it

If we can establish experimental results for our lab mice, we may be able to use the fact that they are mice to make valid inferences to the population of mice. But to the extent to which we can infer from the experimental organism to other possible organisms depends not only on material

qualities but also on the extent to which we can label our experimental subject, situation, and manipulation somehow typical of those others in the world. (Morgan 2003, p. 227-8)

Morgan uses the terms 'sameness' and 'similarity' to refer to the ontological resemblance of experimental models. In the example mentioned above, the results established with the mice are immediately valid to the population of mice. The degree of artificiality of these model organisms that their natural counterparts lack (insofar as lab mice have been manipulated for research purposes) is not problematic because "they are of the same type and the same stuff". That is, they are *representatives* of their population and this is an attribute that validates a range of relatively unproblematic inferences beyond the particular strain but within the species. The extrapolation of results to organisms from other species requires extra care. In this case, a similarity relation must be demonstrated between the experimental subject and the object of interest. On Morgan's view, it must be demonstrated that the former is an adequate *representative* for the latter. In short, "real experiments may allow inference to the *same* kinds of things in the world if they can be considered representative of them and to *similar* things if they are representative for them" (2003, p. 230-1, emphasis added). In both cases, ontological proximity between the experimental object and the reference object is the source of epistemic superiority of laboratory experiments relative to other kinds of experiments.

Mathematical model experiments, on the other hand, can only aspire to be *representations* of processes or objects or relations existing in the world. This ontological difference "creates an immediate and deeper inferential gap between the object of experiment and the object of reference, a gap that is independent of other qualities of the model, for the experimental and reference object are no longer made of the same stuff" (op. cit., p. 230). Moreover, the evaluation of the adequacy of the representation is a more complicated matter because "we have less agreement on procedures for deciding when a representation is a good one" (op. cit., p. 231).¹⁵⁵

In sum, whereas the extrapolation of results from the world of models to the domain of the real world invokes representation of 'otherness', i.e. other kinds of things in the world, the extrapolation of results from the experimental world invokes 'sameness' or 'similarity'. The 'sameness' or 'similarity' between the object of scrutiny and the object of inference conveys a higher degree of confidence to the inferences made.

Morgan expresses this difference yet in another way. Whereas laboratory experiments involve 'replication', model experiments involve 'representation':

In the case of laboratory science, successful experiments depend on accurate *replication* in the laboratory of the elements, changes and outcomes in that part of the world relevant to the question. Only then can inferences be made from the results of experiment to teach us something about the world. In the case of modelling science, successful experiments depend on accurate *representation* in the model of the parts of the world relevant to our questions. Otherwise, we can only learn about the world in the model, which may be rewarding for theory development but not for learning about the world. (Morgan 2002, p. 57 emphasis added)

Thus, whereas experimenters can reproduce the phenomenon of interest in the laboratory (which involves same or similar inputs, processes, and outcomes), modellers can only represent them by having recourse to different media.¹⁵⁶ In economics, the issue is cast in the following terms:

In experimental economics, the validity of experimental results is defended by referring to the design of experiment. Control is dependent on the choice of experimental set up, circumstances and procedures...These choices are guided by the experimenter's need to design the experiment in such a way that economic behavior is made manifest in the experiment. *Although there are arguments as to whether experimental subjects... really do behave "naturally" in the artificial situation* like, for example, managers in industries, nevertheless, they share the quality of being humans, and so that part of the inference gap is surely less than for mathematical model experiments. These design features are then adduced by experimentalists as reasons why the results they find in their controlled situations carry over to the uncontrolled world. This may make it possible to infer to the same type of situations (in terms of the objects, structure and circumstances) in the world, but, *that very same tightness of controls and high levels of specificity involved in the laboratory experimental set-up makes inferences to related or similar situations in the world more problematic.* (Morgan 2002, p. 54-5 emphasis mine)

Though economic laboratory experiments possess the same relative epistemic virtues as the experiments of natural sciences, the derivation of inferences from the former seems to require more caution. The reason is that the 'material' of economics experiments is not quite the same; experimental subjects are not quite made of the same 'stuff' of economic agents in the real world (recall the simplicity and the artificiality criticisms discussed in Ch. 6). Nonetheless, the inferential gap is lower from laboratory experiments than from mathematical models. The most serious problems concern the 'tightness of controls' and the 'high level of specificity' of the experimental systems that render inferences from the experimental world very specific in character and thereby of limited applicability. In contrast, the level of abstraction and simplification of economic events in mathematical model experiments allow inferences to a vast number of circumstances in the economic world that share common aspects with the model situation. However, these are "general inferences" that have little import to specific cases (Morgan 2002, p. 53-7).

Morgan also identifies the tension intrinsic to economics experiments already noted here. On the one hand, the epistemic value of experimental results depends on the participation of experimental subjects because it carries a high potential for the generation of new knowledge. On the other hand, this participation must be controlled so that the knowledge generated can be understood. Consequently, inference from economic experiments cannot be but a local affair.

Morgan's view on experiments is not unique. Francesco Guala also focuses on the role of experiments in providing knowledge about concrete real world situations or

¹⁵⁵ This contrasts with Mäki's terminology in that on his view models are always representatives and representations of something for given purposes.

¹⁵⁶ The term replication is here used in a broad sense to encompass the possibility of reproducing representative phenomena of (for) real world phenomena.

on the problem of external validity.¹⁵⁷ Borrowing from Morrison and Morgan's (1999) conceptualisation of models as mediators, Guala qualifies experiments as 'epistemic mediators' that "help bridging the gap between a theory and its target domain of application", where the target is the non-laboratory phenomenon that scientists seek to explain. On this account, an economic experiment aims at explaining real world phenomena, or more specifically, the postulation of causes, causal processes and causal mechanisms. Guala also argues that experiments are superior epistemic mediators because the relationship between the experimental and the target system "holds at a 'deep', 'material' level since "the same 'material' causes as those in the target system are at work". This contrasts with other mediating devices, such as simulations where the similarity between the simulating model and the target system is "abstract" and "formal" (Guala 2002a, p. 66-7).

Guala is, however, more explicit about what it takes to extend inferences from the experimental to the target system. The mere 'replication' of 'elements, changes and outcomes' does not suffice to justify the external validity of experimental results. The problem of 'causal underdetermination', i.e. the possibility that the same data may be generated by different causal processes, requires an extra step: "one has to show that the system constructed in the laboratory is the same as the one at work in the real world... and that the similarity between 'artificial' results and 'real' phenomena is not illusory" (1998, p. 906). According to Guala this means three things. First, it must be shown that the materials used "resemble as closely as possible those of which the parts of the target system are made". Second, it must be shown that "the components of the mediating device are put together just like those of the target". And, third, it must be demonstrated "that nothing else is interfering" in both the experimental and target systems (Guala 2002a, p. 70). And this can only be established by further empirical testing.

Because an external validity claim amounts to the assertion that the observed features in both the target and the experimental system are generated by similar causal mechanisms, testing this assertion boils down to investigating whether "the probability of observing such correspondence (were the hypothesis false) is low if we have eliminated alternative reasons why such a correspondence might occur, other than the causal similarity between the two systems" (Guala 2003, p. 1202). Again, what follows from this is that inference from laboratory experimental systems is local because it demands a strong similarity between the experimental and the local system. This is why "exporting the lab" allows the safest 'external validity' claims, this is the case where reality is moulded to fit the experimental environment (*op. cit.*, p. 1204, see Ch. 6).

8.5 THE INTERMEDIATE MODE OF INFERENCE

Even though Morgan and Guala are specifically addressing the issue of how and to what extent experimental results may be transferred to concrete real world situations,

and how they differ from those produced by other methods, they provide an unjustifiably limited account of economic experiments. This limitation is intensified in the contrast Guala draws between economics and other 'laboratory sciences'.

Following Hacking's characterisation of laboratory sciences (Hacking 1992), Guala depicts economics as a non-laboratory science "whose claims to truth do *not* answer primarily to work done in the laboratory". Its primary aim is instead "to explain and control non-laboratory phenomena" (Guala 1998, p. 901 emphasis in original). This conception possibly conveys the goal experimental economists should aspire to. The problem is that it leaves unaccounted a significant bulk of experimental practice that falls short of that final and ultimate goal. As the present work testifies, most of economics experiments are either inspired by theory or by previous experimental results. Only a minority of experiments have attempted to explain naturally occurring phenomena. Thus, the majority of experimental economics (and theoretical economics inspired by it) does revolve around the phenomena created in the laboratory. Nonetheless, the experimental phenomena and the theoretical accounts they have inspired can be presumed valid scientific work.

The conception of an experiment as an epistemic mediator not only renders the justification of experiment dependent on some pre-defined target system, but it also leads to an excessive concentration on 'external validity' issues. The social epistemology put forward here obviates the need of having the epistemic value of experiments dependent on some predefined and concrete real-world referent. It does that by providing grounds for an intermediate level of inference that can be placed in-between the 'internal' and 'external' kinds of inference.

The claim I am making here is that *an economic experiment provides knowledge about the human and the social worlds prior to the correspondence that may be established with a concrete real world situation*. This possibility is afforded by the ontological similarity between experimental systems and real world systems. The experiments reviewed in chapter 7 show just that. The ultimatum game experiments did not have as a reference object a particular phenomenon of the real world, a theoretical object of reference, viz. game theory, guided them instead. Rather than reproducing relevant aspects of a target from the real world, these experiments created experimental situations that could feedback into economic theory. Nonetheless, it is undeniable that the results of these experiments convey information about human preferences and behaviour and how these may be affected by the context of interaction (*cf.* paragraphs in italic at the end of section 7.5). The market experiments, too, did teach how institutions affect aggregate outcomes prior to their application to real world markets. As explained above, the ultimatum experiments *qua* behavioural experiments provide knowledge about human behaviour; the market experiments *qua* technological experiments provide knowledge about how to control human behaviour for specific purposes. What grounds these inferences is the 'human agency' of experimental subjects (in behavioural experiments) or the success in controlling it for intended purposes (in technological experiments). Of course, the question of whether these results apply to a particular situation remains an empirically relevant question. In sum, and recovering the notions of internal and external validity, the point can be synthesised as follows:

¹⁵⁷ As mentioned earlier (Ch. 6), 'internal validity' pertains to the establishment of experimental results in the laboratory and 'external validity' to the establishment of the applicability of these results to non-laboratory contexts.

Whereas internal validity refers to the establishment of valid conclusions to the laboratory world, external validity investigates whether or not these conclusions apply also to concrete real world situations. An intermediate genre of inference is available that extends beyond the inferences made to the experimental world but which does not require appraising the similarity with a predetermined and specific target system. The validity of intermediate conclusions is conferred to by the ontological similarity between experimental systems and the real world. Experimental systems are socioeconomic contexts from which knowledge can be acquired about individual preferences, motivations and behaviour, as well as about socioeconomic institutions. Whether or not a parallel context exists in the real world is a relevant question that can be answered by further empirical work.

An equivalent genre of intermediate inference can be found in economic modelling. In fact, this is the most common inference in economics, which has already been argued for and therefore can be brought to bear on the experimental economics. This is the point to which I now turn to.

8.6 MODELLING AND THE INTERMEDIATE MODE OF INFERENCE

Much of modelling in economics aims at generating inferences from the world of models to the real world. Indeed, it would be a very unsatisfactory justification of model building (both epistemic and social) the derivation of hypotheses that are only true of the model world. Instead, economists strive for the generation of general hypotheses that can be expected to apply elsewhere. Allan F. Gibbard and Hal R. Varian (1978) claim that economic models are and should be more than “exercises in mathematics or apologies for *laissez faire* capitalism” (p. 665). They should provide understanding of the world, instead. As they put it, “[o]ften when a model is presented, only the briefest suggestive remarks are made about its bearing in the world, and yet it seems clear that, when an economist investigates a model, it is often because he thinks the model will help to explain something about the world” (p. 676). Of course, the explanatory power of models also hinges upon the model’s representational attributes (approximate or as a caricature according to the authors), which afford the transportability of the results of the model to real world situations. However, the relation of similarity between the model and its domain of application is not univocal. It is instead a relation between the model and general circumstances common to a variety of situations. On this account, then, explanation consists in ‘casual’ application of the model’s conclusions which provides understanding of “aspects of the world that can be noticed or conjectured without explicit techniques of measurement” (p. 672).

More recently, Sugden (2002) claims that the possibility of inferring from a model to the real world does not require that the former be built upon a concrete real world situation. What justifies inference is instead the extent to which the model is a ‘credible counterfactual world’ that describes how the world could be. The upshot is that the potential ‘reality’ of the entities, properties, relations, and so forth, described in the model worlds bestows credibility to the generalisations obtained from them so that they can legitimately be assumed to be applicable to the real world were it to be arranged in the same way. For instance, when describing Akerlof’s lemons model Sugden says:

The world of his model is much more uniform and regular than the real world, but Akerlof clearly wants us to think that there could be a used-car market which was like his model. The “cars” and “traders” of his model are not just primitives in a formal deductive system. They are, I suggest, cars which are *like* real cars, and traders which are *like* real traders, inhabiting a world which Akerlof has imagined, but which is sufficiently close to the real world that we can imagine its being real. (op. cit., p. 131 emphasis in original)

The high degree of confidence economists have on theoretical models and their scepticism about economic experiments is therefore perplexing. Just as modelling results potentially apply to real world situations without having a concrete target in mind, experimental results potentially apply to non-laboratory circumstances, too. Furthermore, the ontological superiority of experiments – the fact that the experiments are made of the ‘same stuff’ – adds epistemic value to the inferences derived from them.

This is not to deny that the possibility of generating inferences from experimental systems to a set of circumstances in the real world hinge on relations of similarity between experimental entities, properties, and so forth, and those to which they refer to in the real world. As noted above, inference depends on the representational attributes of the experimental systems. The point is simply that the ontological similarity between experimental and non-experimental situations allows this intermediate kind of inference. It is this similarity that justifies experimenters’ ease at constructing experimental worlds irrespective of whether a real world referent exists. It is this similarity that allows extending experimental conclusions beyond the reality of the experimental world.

The neglect of intermediate inference by both experimenters and the students of experiments is not surprising considering the hostility towards experimental economics. Economists had first to argue for the ‘external validity’ of economic experiments. After all, the ultimate goal of laboratory economics is to generate information about concrete real-world situations. Therefore, they had to demonstrate that experimental results could produce knowledge about real aspects of the world outside the laboratory.

The making of inferences of a general kind is bold. However, the possibility of intermediate inferences is supported by the internal and external validity of economic experiments. On the one hand, simple, isolated and controlled experimental microeconomies afford the derivation of clear-cut conclusions. On the other hand, the fact that some of those conclusions apply to concrete situations provides reasons to expect their wider applicability for a class of similar situations. Nonetheless, the intermediate kind of inference differs from the internal and external modes. Internally valid inferences are claims about laboratory worlds. Externally valid inferences rely on the possibility of creating (in both the lab and in the field) highly structured environments that severely constrain human behaviour for intended purposes. Intermediate valid inferences are general conclusions about individual preferences, motivations and behaviour, or socioeconomic outcomes that can be observed in the laboratory and in real world economies.

8.7 ILLUSTRATING INTERMEDIATE INFERENCE

Establishing internal and external validity is more difficult with behavioural experiments than with technological experiments. The higher level of 'human agency' of behavioural experiments requires a more extensive series of experiments to interpret experimental results. The high level of 'human agency' of potentially parallel situations in the real world, in turn, makes it virtually impossible to ascertain whether or not external validity holds in there. The lack of control of naturally occurring process makes it extremely difficult to interpret the motivations underlying human behaviour. As we have seen, the almost exclusive possibility of exercising control in the laboratory is where the interest of conducting experiments lies. Nonetheless, behavioural experiments allow generating general inferences about human behaviour which may be amenable to a wide range of 'casual applications', just as model inferences are.

In this regard, it is worth mentioning Ernst Fehr and Simon Gächter's (2000) 'economics of reciprocity' which accounts for what is deemed to be a general pattern of behaviour. The economics of reciprocity is built upon experimental phenomena produced by the ultimatum and public good experiments. As illustrated in the previous chapter, experimental data has produced robust evidence for 'reciprocity'. That is, for the human propensity to respond positively (negatively) to friendly (hostile) actions even when there are no material gains in view. However, it may be difficult to ascertain whether a concrete real world situation constitutes an instance of reciprocity. For it is virtually impossible to ascertain by non-experimental means what were the motivations of individuals in particular circumstances. However, what has been learnt about reciprocal behaviour in experimental contexts offers some understanding of real world situations. The authors take experimental evidence as relevant to provide understanding of recent events in Kosovo when many Albanian refugees took revenge after the victory of NATO over Serbian forces, or the rise in employees' theft rates after firms have cut employees' wages, or the social ostracism exercised by coworkers against strike breakers during and after industrial disputes, among others (op. cit., p. 161).

Even though none of the ultimatum experiments had in mind war settings or labour market conflicts, they do improve our understanding of them irrespective of the possibility of determining the actual factors in operation in those particular instances. In the examples given above it is impossible to determine the actual motivations of reciprocity. But the already extensive array of laboratory experiments suggests that real world behaviour may be explained by the intention to punish (reward) behaviour deemed unfair (fair). To conclude, behavioural experiments suggest that intermediate inferences provide understanding of situations in which it is difficult to scrutinise the actual causal factors in operation. However, intermediate inference may be improved by additional empirical evidence from ethnographic, historical, psychological and sociological studies as well as from mere everyday observation. In contrast, technological experiments provide understanding of real microeconomic institutions and the causal factors in operation therein.

8.8 SUMMARY

In this chapter I argued that the human and the social constitution of economic experiments render the experimental method a superior device from which inferences can be made with a higher degree of confidence. This superiority allows the making of an intermediate mode of inference that has not yet been addressed by the students of experiment. The intermediate mode of inference produces general conclusions that can offer understanding of real world situations. This kind of inference, however, has a different content according to the kind of experiment. While technological experiments supply inferences about microeconomic institutions, behavioural experiments provide inferences about human motivations and behaviour. The high level of 'human agency' of behavioural experiments makes it difficult the application of this general kind of inference to particular instances. In contrast, technological experiments produce results that apply to particular real world situations. In both cases, the ability to generate meaningful knowledge derives from the human and social constitution of the experimental systems of economics. The fact that the bulk of experimental practice revolves around understanding laboratory phenomena is not as problematic as ongoing discussions in the field suggest. The human and socioeconomic ontology of experimental systems relieves the requirement of demonstrating tight similarities before inferences be made. It suffices that experimenters learn something relevant about the reality created in the laboratory.

CHAPTER 9

PREFERENCE REVERSALS OR GOOD EXPERIMENTAL PRACTICE IN ECONOMICS

9.1 INTRODUCTION

The preference reversal research program is one of the best well-studied episodes of experimental practice and it is often used to illustrate the fragility of experimentation in economics (Tammi 1999a; Starmer 1999a, 1999b; Siakantaris 2000), and the negligent (Camerer 1995) or dogmatic (Hausman 1992) attitude of economists when confronted with experimental evidence. In this chapter I show that, in fact, the preference reversal program is an example of good scientific practice that illustrates the epistemic attributes of experimentation in economics. In particular, it highlights the social dimension of the knowledge production process of experimental economics. It shows how experimental knowledge is produced in an incremental way by series of experiments intended to pose and respond to critical charges levelled against previous results. And how the incremental way by which knowledge is produced depends on the scientific community's receptiveness and responsiveness to criticism. Finally, it highlights the importance of interdisciplinary dialogue for identifying the most ingrained beliefs and thereby the removal of obstacles to scientific progress. The chapter thereby rounds up the main results of SEEE which are brought to bear on the preference reversal research program.

The chapter is organised as follows. Section 2 presents the preference reversal phenomenon produced by the psychologists. Section 3 describes economists' reaction to this anomaly to Expected Utility Theory. Section 4 identifies the conditions that had to be met for economists to accept the relevance of the phenomenon to economics. Sections 5 and 6 present the alternative explanations put forward to explain it and section 7 discriminates among them. The significance of the preference reversal phenomenon to economic theory is scrutinised in section 8. Section 9 analyses economists' recent attempts to lessen its impact to economics and section 10 presents an attempt at organising the conflicting views in the field. Section 11 concludes the chapter.

9.2 THE PREFERENCE REVERSAL PHENOMENON

The preference reversal (PR) phenomenon, first reported and reproduced by psychologists (Lichtenstein and Slovic 1971, 1973; Lindman 1971; Slovic 1975), consists of an inversion of preferences revealed in two tasks: a choice task and a

valuation task.¹⁵⁸ In the choice task preferences are expressed in terms of a choice between two lotteries, conventionally known as the probability bet (P bet) with a very high probability of winning a modest amount of money and the money bet (\$ bet) with a low probability of winning a relative large amount of money. In the valuation task preferences are expressed in terms of the prices individuals assign to these lotteries. These prices presumably denote the reservation prices of the lotteries, that is, the maximum amount the individual would be willing to pay for the lotteries, or the minimum price they would be willing to sell them. The PR phenomenon consists of having a large proportion of subjects choosing the P bet in the choice task but assigning a higher price for the \$ bet in the valuation task. The opposite reversal is, however, less common. Subjects rarely choose the \$ bet and price higher the P bet.

In fact, the PR experiment was designed to test the conjecture that preferences depend on the decision problems and on how these are presented to the individuals. In particular, it was designed to test the hypothesis that choosing and pricing are decision problems that appeal to different cognitive processes. The expectation was that choosing is correlated with probabilities whereas pricing is correlated with the monetary prizes of the bets. If correct, it would yield the result that subjects tend to choose the P bet while placing higher prices for the \$ bet, which was indeed observed.

The PR phenomenon caught the interest of economists because it represented a challenge to Expected Utility Theory (EUT) which predicts that rational agents place higher reservation prices on the objects of their choice.¹⁵⁹ This means that the PR phenomenon represented a serious incoherence between experimental evidence and a basic theoretical prediction of a well-established theory of economics. Inasmuch as this incoherence was generated by an item of a distinct scientific culture, i.e. that of experimental psychology, the economists' immediate response was to undermine the significance of the PR phenomenon to economics. In this vein, David Grether and Charles Plott (1979), the first economists to study the phenomenon, thought that the PR would not be reproduced by the experimental systems of economics. But the PR phenomenon resisted the experimenters' expectations. It was observed even in an experiment designed to eradicate it. Not so surprisingly, the experiment did not convince the profession about the significance of the phenomenon to economics.

Economists remained sceptical and, consequently, further attempts were made to eliminate the occurrence of the phenomenon. Because they continued to fail in doing so, they then tried to belittle the impact of the phenomenon to economics. Three strategies were pursued as a reaction to the PR phenomenon. The first strategy, following on Grether's and Plott's approach, tried to undermine the impact of the phenomenon by invoking the methodological differences between experimental economics and experimental psychology (Reilly 1982; Pommerehne *et al.* 1982). The second strategy argued that the PR phenomenon was an artefact of the experimental systems. In particular, it attributed the reversals to the procedures used to elicit preferences (Holt 1986; Karni and Safra 1987; Segal 1988). The third strategy, minimised the impact of the reversals by theoretical means while acknowledging the existence of reversals. To this end, alternative theories were put forward that amended the less crucial parts of EUT while leaving its fundamentals intact (Loomes and Sugden 1983). More recently, the first strategy was taken up again. However, this time it opened the debate about the scope of the discipline of economics (Harrison 1994; Bohm 1994, Cox and Grether 1996).

The fact that after three decades of research the debate still revolved around the same sort of questions was interpreted as a serious scientific failure and, in particular, as an indication of a shortcoming of the experimental method. The PR research program was, in effect, presented as an instance of the Duhem-Quine thesis (see section 6.3). Because in experimental economics the crucial set of auxiliary assumptions consist of the four precepts of experimental economics (see sections 5.2 and 6.2), the Duhem-Quine problem often pertains to the difficulty in establishing whether or not the four precepts of experimental economics are verified. As a result, "[i]t must always be a matter of judgment and, more pointedly, a matter of judgment about controversial matters such as whether the incentives were big enough, whether the subjects fully understood the design, and so on" (Starmer 1999a, p. 18). In addition, it is also possible to protect theories from refuting experimental evidence by restricting their domain of application. In sum, "[t]he data will not speak for itself, and, ultimately, conclusions on these matters will turn on some degree of judgment leaving room for inter-subjective disagreement" (op. cit., p. 19). Any conclusion may thus be challenged and consensus may never be forthcoming. On Starmer's reading, this has been the case with the PR research program. To some participants of the controversy, the experimental evidence is sound and supports the view that EUT is refuted (Camerer 1995) whilst for others the evidence is not reliable because it fails to satisfy the precepts of experimental economics (Harrison 1994). In addition, it has also been argued that PR do not occur where economic theory purportedly applies. That is, in contexts where decisions are consequential and where there are opportunities for learning with experience (Binmore 1994). Starmer concludes that after three decades of extensive research too much controversy continues to abound and as yet no question has been conclusively answered (1999a, p. 24).

Timo Tammi also takes the PR dispute as an instance of the Duhem-Quine problem, but one that is understood in the light of the cost-benefit analysis inherent to the scientific process of social negotiation:

¹⁵⁸ This category of experiments is distinct from the others reviewed so far in that they aim to elicit subjects' preferences rather than to induce them.

¹⁵⁹ Expected utility theory posits that rational action is the action that yields the highest expected utility, where expected utility is the sum of the utilities of each possible outcome of an action multiplied by the probability of the outcome's occurrence ($\sum u_i \cdot p_i$). If preferences obey a particular set of axioms (completeness, transitivity, continuity, reduction, independence, monotonicity) then EUT yields predictions about the behaviour of a rational agent. The completeness axiom specifies that all objects over which individuals have preferences can be compared or ordered; the transitivity axiom specifies that if an individual prefer object a to b, b to c, then he also prefers a to c; the reduction of compound lotteries axiom asserts that an individual is indifferent between a compound lottery that has another lottery as one of its prizes and this lottery; the continuity axiom states that if an individual prefer a to b, and b is preferred to c, then there will be probabilities such that he will be indifferent between b and a lottery with a and c as prizes; the independence axiom says that if an individual is indifferent between two prizes, a and b, then he is also indifferent between two similar lotteries in which one has a as prize and the other b; monotonicity says that between any two lotteries involving only a most and a least preferred alternative, the individual prefers the lottery where the least preferred alternative is less probable.

The core idea is that in a situation where there is a set of inconsistent statements, a major aim of each scientist is to alter the composition of the set so that the inconsistency disappears. Which one of the available ways is presented and advocated by one participant to resolve the inconsistency depends on his/her cost function, that is, the net costs resulting from the fulfilment of a proposed resolution. As a general principle, smaller costs are preferred to greater costs. (1999a, p. 267-8)

According to Tammi the costs involved in scientific negotiation are diverse. They include the effect of given solutions to scientists' reputation, the protection of one's scientific human capital, the cognitive costs associated with engaging in new lines of research, the financial costs of research, the costs of rejecting established practices, and so forth. But like Starmer, Tammi also thinks that a final resolution of the dispute is yet to be expected.

Similarly, Colin Camerer (1995) considers that economists have devoted too much time to 'destructive' rather than 'constructive' contributions to the analysis of anomalies of individual decision-making theory:

Destructive tests, often motivated by skepticism, are designed to check whether apparent anomalies are replicable, robust across settings, or might be due to flaws in experimental design. My opinion is that *some* occasional tests of this sort are essential, but too much energy has been devoted to destructive testing with very little payoff. Not a single major recent (post-1970) anomaly has been "destroyed" by hostile replication of this sort. (1995, p. 674 emphasis in original)

Even though controversy abounds, and despite the fact that the PR research program has not provided answers to all of the questions raised, it is nonetheless an example of good scientific practice that illustrates the epistemic virtues of the experimental method of economics. It is now commonly acknowledged that the reversal of preferences is a well-understood phenomenon that is most likely to be observed in contexts where individuals do not have the opportunity to learn from experience. This suggests that EUT may nonetheless preserve its predictive power in structured contexts where feedback and opportunities for learning are available. But its most significant contribution is now in progress: the development of a new research program that takes as problem-situation the study of the processes of individual decision-making.

The view that experiments produce evidence that can support any scientific bias derives from two misconceptions about experimentation. First, it expresses the idea that the exclusive function of the experimental method is the empirical testing of theories which ought to be delivered by a single experiment. Second, it expresses the idea that unanimity is a feasible and a necessarily desirable collective epistemic state. SEEE dismantles such misconceptions. The recognition that an experiment is a device purportedly designed and implemented for providing answers to specific questions implies that more fundamental questions can only be answered by additional experimental and non-experimental research. As will become apparent below, SEEE shows that rather than producing conflicting answers to the same question, the experimental results indeed shed light on various issues. The issues that were not addressed were simply not constituted as problem-situations. To put it bluntly, the alleged controversy is unjustified and the issues left unsolved are the

result of a lack of interest on the part of the economics community at large rather than of problems pertaining to the experimental method.

9.3 THE PREFERENCE REVERSAL PHENOMENON AND THE METHODOLOGICAL DEBATE

The preference reversal research program was brought to economics by the experimenters David Grether and Charles Plott (1979) who were intrigued by the challenge the phenomenon presented to EUT. This is how the economists perceived the implications of the PRs to economics:

A body of data and theory has been developed within psychology which should be of interest to economists. Taken at face value the data are simply inconsistent with preference theory and have broad implications about research priorities within economics. The inconsistency is deeper than the mere lack of transitivity or even stochastic transitivity. It suggests that no optimization principles of any sort lie behind even the simplest of human choices and that the uniformities in human choice behavior which lie behind market behavior may result of principles which are of a completely different sort from those generally accepted. (1979, p. 623)

Even though the phenomenon produced by the psychologists was interesting in its own terms, it could not be taken at face value. The recognition that "psychologists have uncovered a systematic and interesting aspect of human choice behavior" begged the question of "whether this behavior should be of interest to economists". In particular, it raised two questions: "1) Does the phenomenon exist in situations where economic theory is generally applied? 2) Can the phenomenon be explained by applying standard economic theory or some immediate variant thereof?" (p. 624). The answers to these questions were to be found by experimental means. The expectation was that the phenomenon would not be reproduced by the experimental method of economics. The difference between the experimental practices of economics and those of psychology was then the primary explanation of the phenomenon. The experimental method of psychology rendered the 'instrumental model' of the PR experiment inadequate to generate phenomenon of interest to economics. The experimenters were confident that the PR would not occur if the original experiment were to be implemented with the precepts of experimental economics. This was the economists' 'instrumental model'. It was a natural step to take. On the one hand, the principles and practices of experimental economics were different from experimental psychology. On the other hand, EUT was a well-established resource of practice that would not be easily questioned by a result produced by another field of research.

The significance of the PR phenomenon to economics was not obvious because experimental psychology relies on what are considered to be fundamentally different principles.¹⁶⁰ The major distinction between experimental economics and experimental psychology concerns the generalised adoption of financial incentives based on subjects' performance in the former and its exceptional use in the latter. From the viewpoint of economists this methodological difference casts serious

doubts on the ability of psychologists in inducing motivated behaviour. More specifically, in economics the subjects' monetary pay-off depends on their performance (the saliency precept), while in psychology the subjects' behaviour is often inconsequential. That is, in psychology, subjects are asked to take decisions on the basis of hypothetical consequences that are not experienced by them in the course of the experiment. At best, subjects receive a flat fee for their participation. More often they get course credits in their stand. This does not mean that a performance-based reward structure is prohibited in psychology. It is simply not part of widespread practice as it is in economics. But, in fact, some of the PR experiments in psychology did adopt it.¹⁶¹

Even though in the seminal methodological papers of experimental economics (see Chs. 5 and 6) the inclusion of monetary incentives is presented as a sufficient but not necessary procedure to induce relevant and intelligible behaviour, Grether and Plott take monetary incentives as a necessary condition for inducing 'economic behaviour' in the sense that it bestows economic significance to the experimental decisions. They say:

Almost all economic theory is applied to situations where the agent choosing is seriously concerned or is at least choosing from among options that in some sense matter. No attempt is made to expand the theory to cover choices from options which yield consequences of no importance. Theories about decision-making costs do suggest that unmotivated choice behavior may be very different from highly motivated choice behavior, but the differences remain essentially unexplored. Thus, the results of experiments where subjects may be bored, playing games, or otherwise not motivated, present no immediate challenges to theory. (1979, p. 624)

The adoption of a performance-based financial incentive structure was then the first substantial amendment introduced to the 'material apparatus' with the goal of creating an economically significant context. The second major modification attempted to overcome the widely used practice of deception in psychology. The fear was that the practice of deceiving subjects could have generated a reputation among

the students of psychology that experimenters conceal or even mislead subjects as to what the true objectives of the experiment are. If so, the results of PR experiment would be inconclusive. Rather than performing the relevant experimental task, the subjects were responding to whatever they thought the goal of the experimental task was.¹⁶² In order to avoid this effect, the experiment was performed with economics and political science students instead of psychology students who were deemed to belong to a 'very special population'.

The new 'instrumental model' was then implemented on students recruited from economics and political science classes who were told that they were to take part in an economics experiment. The subjects were informed that they were to participate in a decision-making experiment that involved one or more bets for which they had an initial credit of \$7. It was also explained that their reward would be calculated by adding or subtracting to the initial endowment the proceedings of one bet randomly selected at the end of the experiment. The subjects also knew that they would receive no less than \$5.

They then performed two decision tasks. In the choice task, they either picked up the bet they preferred or expressed indifference between them. In the pricing task they indicated the smallest price at which they would be willing to sell the bet.¹⁶³ At the end of the experiment a *random lottery selection procedure* would randomly select one of those bets to determine the subjects' payoff. This means that the rewards of the subjects would be determined at the end of the experiment and would be calculated on the basis of the results of the selected bet. This was the way the experimenters found to control for the effect of wealth changes on individual preferences. To be had, the reversals could not be attributed to income effects.

If the random selection procedure selected a bet that belonged to the choice task, the subject received the result of playing that bet. If the selected bet belonged to the pricing task, the subject received a randomly generated offer price (from a uniform distribution) if this offer price was higher than the price set by the individual. The subject received the result of playing the bet if the offer price was lower. This procedure is called the *Becker-DeGroot-Marschak mechanism* intended to elicit the lotteries' true certainty equivalents.¹⁶⁴ The goal is to have subjects stating their reservation prices, that is, prices for which the individuals are indifferent to either selling the lotteries or playing them out. This mechanism was deemed effective to that end. The reasoning is that subjects have no incentive to place a higher price on a bet which is less preferred for they risk keeping it, neither have they an incentive to place a low price on a preferred bet for they risk losing it for a price inferior to their subjective value.

These procedures were explained to the subjects in the instruction stage in which subjects were explicitly told that it was in their best interest to reveal true reservation prices. In addition, subjects could learn the experimental tasks in a trial prior to the experiment where results could be discussed and questions answered. Under these

¹⁶⁰ The major methodological differences between experimental economics and experimental psychology are well captured in the article by Ralph Hertwig and Andreas Ortmann (2001) in *Behavioral and Brain Sciences* and the commentary that followed, in the same issue, involving both psychologists and economists. This is how Hertwig and Ortmann synthesise those differences: "Whereas economists bring a precisely defined "script" to experiments and have participants enact it, psychologists often do not provide such a script. Economists often repeat experimental trials; psychologists typically do not. Economists almost always pay participants according to clearly defined performance criteria; psychologists usually pay a flat fee or grant a fixed amount of course credit. Economists do not deceive participants; psychologists, particularly in social psychology, often do (p. 384).

¹⁶¹ As explained in chapter 6, in economics a performance-based incentive structure (the saliency precept) is instrumental to control subjects' motivation. Not only is it perceived as a means to render the experimental task relevant to the subjects, but it is also considered necessary to compensate the costs (attention, care in responding, etc.) of the experimental task. The underlying idea is that subjects will devote more effort to the experimental task if they receive adequate retribution in return. Psychologists presuppose that subjects' intrinsic motivation is high and that they usually try to do their best. This means that psychologists consider that effort does not improve substantially with financial incentives and, consequently, that subjects' behaviour is not substantially changed by their introduction. For a comprehensive study on the effect of incentives on human behaviour see Camerer and Hogarth (1999).

¹⁶² For a recent discussion on the use of deception in economics see McDaniel and Starmer (1998).

¹⁶³ The possibility of indifference was another innovation of the experimental design which would avoid counting equally preferred options as an inconsistency.

¹⁶⁴ This procedure, known also as the BDM procedure is named after the authors who proposed it – Gordon M. Becker, Morris H. DeGroot, and Jacob Marschak (1964).

circumstances, or so the experimenters thought, a reversal of preferences could no longer be justified by inadequate motivation, misconceptions about the purpose of the experiment, income effects or indifference. All these were standard experimental procedures in economics (see section 6.2) that would help creating an economically interesting laboratory context.

The implementation of the new ‘instrumental model’ reproduced the standard pattern of the phenomenon. And to the experimenters’ surprise, the inclusion of monetary incentives had a stronger effect on the rate of reversals (67% as opposed to 56% when payment was uncorrelated with decisions). That is, approximately seventy percent of the choices of P bets were inconsistent with the announced selling prices. The opposite reversal occurred for just thirteen percent of the \$ bets. The experimenters also found no significant effect of market terminology nor did they find an important effect associated with the ordering of the tasks (i.e. pricing after choosing and vice-versa). The economists concluded that the PR phenomenon had been reproduced in economically significant circumstances.

It is interesting to note that the experimenters immediately accepted the result that was in conflict with their expectations. Instead of recommending further research to investigate the robustness of the phenomenon under more severe conditions, i.e. where choices could be made more ‘significant’ to subjects, Grether and Plott concluded, without hesitation, that the theory failed in accounting for the controversial phenomenon.

This seems to suggest that the experimenters had a great deal of confidence in the experimental method and, in particular, in two fundamental standards of experimental economics – the use of monetary incentives and the avoidance of deception. This means that the reward structure was deemed adequate to induce the required motivation on undergraduate students. The amounts paid were considered reasonable to reward the tasks the subjects had to perform and, moreover, they were adequately tied to individual performance. Given that the experiment did not pose special problems concerning other-regarding motivations, the reward structure was therefore the main procedure to induce the required motivation on subjects. Nonetheless, care was taken to avoid contact among the subjects. In sum, the four precepts of economics (nonsatiation, saliency, dominance and privacy) were satisfied (see Chs. 5 and 6).¹⁶⁵

Since the experiment reproduced a result by means of what was considered an adequate design and what were considered superior procedures, Grether and Plott accepted it. The three-way coherence was achieved by accepting the new ‘experimental procedure’ and the new ‘instrumental model’ while rejecting the original ‘phenomenal model’ – the hypothesis that preference reversals do not occur in economically significant situations. In other words, the decision to accept that the PR phenomenon occurs in economically significant situations was basically supported by the conceptual understanding of what an adequate experiment in

economics amounts to and the belief that actual practice conformed to it. Nonetheless, the observation of the PR did not warrant the rejection of EUT. This was a unique disconfirmation of a widely accepted and well-established theory:

The fact that preference theory and related theories of optimization are subject to exception does not mean that they should be discarded. No alternative theory currently available appears to be capable of covering the same extremely broad range of phenomena. In a sense the exception is an important discovery, as it stands as an answer to those who would charge that preference theory is circular and/or without empirical content. It also stands as a challenge to theorists who may attempt to modify the theory to account for this exception without simultaneously making the theory vacuous. (Grether and Plott 1979, p. 634)

9.4 THE COLLECTIVE ACCEPTANCE OF THE PR PHENOMENON

The view that the preference reversals phenomenon is a significant phenomenon to economics did not gain immediate collective assent. EUT was a simply a too well-established theory to be undermined by an insufficiently understood phenomenon. Doubts about the significance of the PR phenomenon continued to revolve around the effectiveness of the incentive structure to induce ‘economic behaviour’.

Robert Reilly (1982) doubted that Grether and Plott’s experiment generated a ‘real’ phenomenon. The problem, according to Reilly, was that subjects might have perceived the decision-problem as affecting the income of the experimenter rather than their own wealth. Moreover, he suspected that the experimental tasks could have been too complex to be understood by the subjects. Reilly designed a new experimental set-up intended to raise subjects’ understanding of the experimental task and their motivation. Extra information about the decision tasks was provided. In particular, subjects received the bets’ expected values and were clarified about its meaning. They were also informed that it was irrational to record selling prices in excess of the possible win of the lottery and to record a payout price higher than the amount that could be lost on that gamble. The understanding of the experiment was also enhanced by reducing the number of participants and by allocating more time to the clarification of doubts in the instruction phase. Finally, the running time of the experiment was increased to allow for greater care in responding to the experimental tasks. In sum, Reilly improved the transparency of the experimental situation so that the observed behaviour could be interpreted in the light of the experiment’s ‘material procedure’ rather than on the basis of subjects’ misconceptions.

To induce the idea that it was their own money that was at stake, the subjects received \$9 at the beginning of the experiment \$5 of which would be immediately kept while the remaining would be used in the course of the experiment. Whether they would lose some or all of the remaining \$4 or win additional cash depended on the decisions they were about to make. To render decisions more relevant to the subjects, the bets were modified to entail substantial losses. The BDM procedure was also modified to extend the interval of the uniform distribution from which the offer prices were drawn, which could now generate negative offer prices. Under the new design

¹⁶⁵ It is worth noting that the PR experiments bring to the fore the prohibition of deception as other principle of experimental economics, which had not yet emerged in the previous accounts of experimental practice. This is an indication of the importance of the interdisciplinary dialogue in rendering salient basic presuppositions not always consciously held.

subjects could pay to avoid participating in a lottery with a loss and could place higher prizes so as to raise the chance of playing the bets with high prizes.

The implementation of the new design produced the predicted pattern of reversals, though at a lower rate (34%). Reilly then concluded that the results “provide further confirmation of preference reversal as a persistent behavioral phenomenon in situations where economic theory is generally applied” (p. 582). The reduced rate of reversals in turn raised the question of whether a further strengthening of monetary incentives and the provision of additional information could continue to reduce the rate of preference reversal. If so, it could be argued that “individuals *are* likely to be consistent in making decisions on alternatives that *matter* to them when the principle characteristics of the alternatives are sufficiently comprehended” (ibid, emphasis in original). Reilly thus produced a result that was somehow coherent with his expectations. Insofar as the net effect of the corrections he introduced to the ‘instrumental model’ moved in the predicted direction, he could still venture that the phenomenon could be eliminated and EUT saved.

Werner Pommerehne, Friedrich Schneider and Peter Zweifel (1982) conducted an experiment along the same lines. The Swiss team regarded the amounts that could be won or lost too small to induce rational behaviour. They decided to increase the face value of the stakes of the bets, nominated in Swiss Francs, by a factor of 100. The conversion rate of the Swiss Francs to real cash payoff would be unknown until the end of the experiment. In this experiment, subjects were given Sfr300 and told that the total cash payoff available for all participants was Sfr2,000 which meant that the payoff of each participant depended on their performance relative to that of others. As with Reilly’s experiment, other features of the experiment were altered to study other aspects of the decision-making process, such as the effect of learning and the possibility of keeping track of past experience. But again, they continued to reproduce the PR phenomenon obtaining a forty-five percent of the predicted pattern of reversals. They concluded that “[e]ven when the subjects are exposed to strong incentives for making motivated, rational decisions, the phenomenon does not vanish” (p. 573).

Even though the experiments conducted by Reilly and by Pommerehne *et al.* were designed to correct aspects of the experiment developed by Grether and Plott they did not use it as a benchmark. The Reilly experiment improved both the incentive structure and the intelligibility of the experiment while Pommerehne *et al.* introduced a radically different incentive structure. As Grether and Plott (1982) noted in a commented reply to Pommerehne *et al.*, “[t]he subjects pools differed. The experiments differed. The language differed, so naturally the questions were “framed” differently. Substantially different motivation conditions were imposed” (p. 575). Therefore, the fact that the rate of reversals was lower did not warrant the claim that this reduction was due to a higher incentive structure. In fact, Grether and Plott even doubted that the incentive structure was stronger. Raising the face value of the bets using a monetary unit of which the conversion rate is unknown does not ensure that the incentives are greater.¹⁶⁶

At any rate, they all agreed that the PR phenomenon subsisted attempts at eliminating it. The Grether and Plott’s experiment produced a sixty-seven per cent rate of standard reversals which supported the conclusion that the introduction of monetary rewards does not eradicate the phenomenon. Reilly’s experiment generated a rate of reversals of about thirty-four per cent which showed that increased motivation and a more transparent experimental design could reduce the rate of reversals. The experiment of Pommerehne *et al.* produced a forty-five per cent rate providing evidence that reversals are robust to payoff differentials between bets, learning and registered track of past experience.

Thus, the set composed by these experiments was successful in reproducing the preference reversal phenomenon. Three different experimental systems involving three different experimental apparatuses did produce a reversal of preferences. Or to put it in another way, three different experimental designs involving different incentive structures continued to produce a significant rate of PR. But perhaps most significant is the fact that preference reversals persisted in the face of substantial efforts made to eliminate it. That is, the PR phenomenon represented an actual resistance of the experimental subjects to the economists’ insistent attempts to tame it. The economists’ predisposition to reject the relevance of the preference reversals therefore conveys a stronger epistemic value to their failure in doing so. As explained in chapter 5, the confirmation of expected results is epistemically weaker than the rejection of prior expectations. The rejection of prior expectations is epistemically significant because it represents a resistance of the object of scrutiny when scientists’ social, cognitive and psychological biases pull in the opposite direction. This was exactly the case of these three sets of experiments. All the modifications introduced to the PR experiment were meant to reduce the range of options of the experimental subjects by rendering one course of action more salient – that consistent with EUT.

The fact that there were no conclusive results about the isolated effect of the incentive structure of the various experiments is not a fault to be attributed to the experimental method. It should be instead attributed to the interests of the experimenters who were not concerned with testing the effect of particular aspects of the incentive structure but with the eradication of the PR phenomenon. No doubts, however, remained that the PR phenomenon is a robust result that occurs in ‘economically significant contexts’ where subjects are adequately motivated. If doubts remained they were further reduced by subsequent experiments [such as the arbitrage experiment by Berg, Dickhaut and O’Brien (1985) and the market experiment by Knez and Smith (1987)] and by subsequent attempts at explaining the cause of the reversals. In the words of Graham Loomes:

The main point emerging from the numerous replications and modifications of the basic experimental design is that the phenomenon stubbornly refused to go away. Providing monetary incentives to encourage careful choice and elicit accurate and truthful valuations, controlling for possible wealth effects and making prices quite large – none of these modifications caused the phenomenon to disappear, nor even to be much reduced. Since the phenomenon appeared to be so robust, effort and attention began to focus on providing an explanation. (1991, p. 600)

¹⁶⁶ This criticism is consensual in the field. ‘Play’ money raises more problems than helps solving motivational issues (e.g. Harrison 1994).

9.5 EXPLAINING THE PR PHENOMENON AWAY: THE TECHNICAL ROUND

The PR phenomenon challenged a well-established theory of individual decision-making under risk. It is therefore not surprising that having excluded the possibility of rejecting the relevance of the phenomenon on the basis of methodological divergences with experimental psychology, economists turned their attention to the appraisal of the particular experimental systems. Their expectation was that the reversals could be attributed to an artefact of the experimental procedures rather than a manifestation of inconsistent preferences. Two techniques were the principal targets of scrutiny: the *random lottery selection procedure* (RLS) and the *Becker-DeGroot-Marschak mechanism* (BDM). The main charge was that these procedures do not warrant the elicitation of true selling prices and therefore they lead to the misreporting of the selling prices.¹⁶⁷

Charles Holt (1986) conjectured that the PR phenomenon could be an artefact of the RLS procedure if the independence axiom was not verified.¹⁶⁸ He argued that if subjects perceived the experiment as a two-stage lottery (the first stage consisting in the selection of the bet to be played out and the second stage consisting in the actual playing of the bet of the pricing task) the choice between the P bet and the \$ bet would be a choice between two compound lotteries. If the independence axiom held the choice between the two compound lotteries would be equivalent to the choice between the two bets. But if the independence axiom did not hold then the choice between the two compound lotteries is different from a direct choice between the two bets. On this reading, the subjects could be choosing the compound lottery containing the least preferred bet (the bet with the lowest selling price) but they would actually rejected it if they faced a direct choice between the two. In short, the compound lottery would not reveal the subjects' preferences over the bets.

It should be noted that this is a mere conjecture. Holt acknowledges that there could be other explanations for the reversals. The violation of the independence axiom is simply the most convenient. This is the case because it is the least damaging to EUT and, in addition, there were at the time alternative theories that made do without it:

Preference reversals could also be generated by intransitivities, but to abandon transitivity would be a drastic step that would make it difficult to construct a formal choice theory with empirical content. The transitivity assumption is needed for the existence of a utility functional that represents preferences over lotteries; independence is a strong assumption about the functional form of this utility functional (that pertains to the linearity with respect to the probabilities). There has been a considerable amount of recent work on preference theories that involve weaker versions of the independence axiom and that permit more general functional forms. (Holt 1986, p. 514)

Edi Karni and Zvi Safra (1987) also consider that the preference reversals may be caused by a violation of the independence axiom. But to these authors the problem lies on the ability of the BDM procedure in eliciting the certainty equivalents of the lotteries. The economists demonstrate this by showing that within the context of theories of choice under risk¹⁶⁹ there is a class of models, which do not follow the independence axiom but verify the other axioms of EUT (completeness, transitivity, continuity and monotonicity), for which the maximisation of the value of the compound lottery does not require the elicitation of the lottery's certainty equivalents. That is, for this class of theories maximising agents may reveal a preference ordering between two compound lotteries consistent with an opposite ranking of selling prices. However, these results obtain only under very strict specifications of the model, in particular those pertaining to the utility function and the probability transformation function, and for the restrict class of lotteries used in the demonstration. In short, Karni and Safra only show that some of the reversals might be explained by a violation of independence rather than transitivity.¹⁷⁰

These interpretations were mere conjectures that hypothesised that the observed behaviour might not have been a manifestation of a reversal of preferences. The wide range of possible interpretations suggests that the PR experiments were not very transparent systems. The PR phenomenon could either be interpreted as a behavioural anomaly or as the outcome of a faulty material procedure. However, the results of the PR experiments could be rendered more transparent by the design and implementation of more experiments.

Noting that none of the interpretations put forward could discriminate between the different causes of reversals, James Cox and Seth Epstein (1989) conducted a new experiment to that effect. To exclude the attribution of reversals to violations of the independence axiom (or the reduction axiom),¹⁷¹ the experiment avoided the RLS procedure and the BDM mechanism. Two major modifications were introduced to the material apparatus. First, subjects would be paid for each decision rather than for the decision selected randomly at the end of the experiment. Second, the elicitation of the lotteries' certainty equivalents was avoided. Because the ranking of preferences would suffice to identify the preference reversals, the pricing task did not need to elicit subjects' true selling prices.

In this experiment, subjects would state prices for both lotteries to decide which among the two they would keep. They would keep and play the lottery to which they had given the highest selling price and they would receive a fixed price in return of the lottery to which they had given the lowest selling price. Under these circumstances, subjects would reveal consistent preferences if they placed a higher selling price on the

¹⁶⁹ These include generalised expected utility (Machina 1982); weighted utility theory (Chew and MacCrimmon 1979; Fishburn 1983) and expected utility with rank-dependent probabilities – EURDP (Quiggin 1982, Yaari 1987).

¹⁷⁰ It should be noted that the economists did not refer to previous experiments that did not have recourse to the BDM mechanism (e.g. Lichtenstein and Slovic 1971) and still obtained a high rate of reversals.

¹⁷¹ For economy of space, the attribution of the preference reversal to the violation of the reduction axiom is not reviewed here (cf. Segal 1988). But the strategy is the same. It demonstrates the possibility of chosen compound lottery to be reduced to an actuarially equivalent one-stage lottery that is given a lower price.

¹⁶⁷ For a more detailed account of the 'technical round' see Guala (2000). Guala interprets this debate as an instance of the 'theory ladenness of observation' in that "the phenomenon at stake is inconsistent with expected utility theory, but the instruments used to observe the phenomenon are constructed on the hypothesis that expected utility is correct" (p. 56). See also Guala (2005b, Ch.5).

¹⁶⁸ The independence axiom asserts that if lottery A is preferred to lottery B, then the compound lottery (A, p; C, 1-p) is preferred to (B, p; C, 1-p). The implication of this axiom is that only the outcomes that distinguish two lotteries are relevant to choosing between them.

lottery for which they had shown a preference. Otherwise, they would exhibit a preference reversal or, better, a ‘choice reversal’.¹⁷² This experiment produced a rate of reversals of about thirty-five per cent which could not be attributed to violations of the independence or the reduction axioms. But it failed to obtain the standard pattern of reversals given that the predicted and the unexpected type of reversals were of the same magnitude. The statistical analysis of the data excluded the presence of significant wealth effects or the effect of subjects’ learning experience to the frequency and the pattern of preference reversals, and framing effects. The series of experiments that followed produced further evidence for the hypothesis that the PR could not be fully explained as an artefact of the RLS procedure (Beattie and Loomes 1997; Cubitt *et al.* 1998) or that of the BDM mechanism (Starmer and Sugden 1991, Keller *et al.* 1993).

In sum, the technical round reinforces the argument that experimenters’ practice was epistemically sound. In the presence of a well-established theory and weak refuting evidence it was a natural step to investigate alternative explanations to the PR phenomenon. The extension of the knowledge production process by other economists was not trivial. It generated a strong consensus that the PR phenomenon represents a far more damaging challenge to EUT than that implicated by a violation of the independence or the reduction axiom.

9.6 EXPLAINING THE PR PHENOMENON

The attribution of the PR phenomenon to a violation of the independence and the reduction axioms were the least harmful to EUT while being compatible with the emergent non-expected utility theories that satisfied the two fundamental principles of EUT. Firstly, they all assumed that individuals have clear and well-specified preferences which they apply to any decision problem. Secondly, they assumed that individuals make decisions which are optimal in the light of their preferences.¹⁷³ A bolder interpretation of the PR phenomenon attributed the anomaly to a violation of a more fundamental axiom of EUT: the transitivity axiom. However, the basic principles of EUT remained unchallenged. This was another interpretation of the PR phenomenon by the economists which provides further evidence for the lack of transparency or stringency of the PR experiments.

VIOLATING TRANSITIVITY

Graham Loomes and Robert Sugden (1983) explained the phenomenon of preference reversals within the framework of regret theory which relaxes the axiom of transitivity while being consistent with the basic presupposition that individuals

have stable preferences and that they are rational and optimising agents.¹⁷⁴ The key idea of regret theory is that in contexts of uncertainty individuals anticipate and take into account the sensations of regret and rejoicing that result from the rejection of available options in the past:

The essential notion underlying regret theory is that people tend to compare their actual situations with the ones they would have been in, had they made different choices in the past. If they realize that a different choice would have led to a better outcome, they may experience the painful sensation of regret; if the alternative would have led to a worse outcome, they may experience a pleasurable sensation we call “rejoicing”. (1983, p. 428)

The problem the individual faces in pairwise choices, as in the PR experiments, is formulated as follows. In pairwise choices the individual faces a choice between pairs of ‘actions’ each of which entails a set of state-contingent ‘consequences’. For any given state of the world, the utility the individual derives with consequence x_1 by choosing action A_1 is affected by the utility that he would have derived had a different action A_2 resulted, for the same state of the world, in a different consequence x_2 . If x_2 is a more desirable consequence than x_1 the individual may experience regret, that is, the unpleasant feeling due to the observation that he could have been in a much better situation had he chosen differently. Conversely, if x_1 is a more desirable consequence he may experience rejoicing, that is, the pleasure associated with the confirmation that he took the best decision. The main assertion of regret theory is that the individual chooses the action that maximises rejoicing and/or minimises regret. Since large differences between what individuals get from a chosen action and what might be received from an alternative give rise to large regrets, individuals prefer greater certainty to less (P bet over \$ bet) than they otherwise would.

VIOLATING PROCEDURE INVARIANCE

A more radical interpretation takes the results of PR experiments not as evidence of a violation of some of the axioms of EUT but as evidence against a taken for granted experimental presupposition – the principle of procedural invariance – which takes elicitation procedures to be neutral mechanisms of eliciting preferences. On this view, the preferences revealed in the experiments are to be explained by the experimental tasks themselves which induce a particular pattern of preferences on the subjects. They depend on the way the decision-problems are presented and on the particular heuristics or computational mechanisms they render most salient.

In fact, as mentioned earlier, Lichtenstein and Slovic (1971) produced the PR phenomenon in an experiment purposively designed to demonstrate that the response mode affects the answers to decision problems. Having noted in an early

¹⁷² The phenomenon observed is dubbed a ‘choice reversal’ to distinguish it from the standard reversals produced by a certainty equivalent elicitation mechanism. The reason is that the pricing task is, in effect, a choice task between two bets.

¹⁷³ For these reasons Chris Starmer dubs these recent non-expected utility theories “conventional” choice theories (2000).

¹⁷⁴ Regret theory was originally developed by the authors to account for other violations to EUT, namely the ‘common ratio effect’, the ‘common consequence effect’ or the Allais’ paradox, and the ‘isolation effect’ (see Loomes and Sugden 1982). Other versions of regret theory were independently developed by Bell (1982) and Fishburn (1982) at the same time.

experiment that pricing is primarily determined by the payoffs and that choices between gambles are primarily influenced by the probabilities of winning or losing (Slovic and Lichtenstein 1968), the psychologists conjectured that if people process information differently when making choices and setting prices, then it should be possible to construct pairs of gambles for which the same individual would choose the high probability and low pay-off bet but set a higher price for the low probability and high payoff bet. That is, the PR experiment was purposively designed to investigate whether or not individuals choose the bet for which they set the lower price. Having confirmed this expectation, the experimental process of knowledge production ended. The experimental procedure, interpreted through the instrumental model produced a result that the psychologists could understand. The natural extension of the knowledge production process was then to develop feasible interpretations for the new phenomenon by conceptual means.

Slovic and Lichtenstein (1983) put forward the ‘anchoring and adjustment’ hypothesis according to which decisions are taken by a process of ‘anchoring’ on certain key features of the problem which are then ‘adjusted’ by reference to other elements of the problem. Accordingly, in a choice task subjects first ‘anchor’ on the probabilities of winning and then revise their choice to take into account the differences in payoffs. Conversely, when setting a price on an attractive bet the subjects ‘anchor’ on the payoffs and adjust it downwards to take into account the probability of winning or losing as well as the amount that could be lost. Because the adjustment process is never complete, the ‘anchor’ substantially affects the subjects’ decisions and it is the cause of the different orderings of a given pair of bets. The ‘anchoring’ and ‘adjustment’ hypothesis therefore predicts the choice of the P bet for those pairs with a larger \$ bet loss relative to the P bet loss, and the setting of higher prices for the \$ bets that have a larger win relative to the P bet win.

Other equivalent explanations were proposed such as the *prominence hypothesis* (Tversky *et al.* 1988) according to which decision-making depends on the prominent dimension of the problem at hand. On this view, choice invokes qualitative reasoning and ordinal considerations, whereas pricing requires quantitative assessment and cardinal considerations. Similarly, the *scale compatibility hypothesis* (Tversky *et al.* 1990) conjectured that the weight of a particular feature of a task is enhanced by its compatibility with the response mode. Because incompatible attributes require additional mental transformations, which increase effort and error, they have a smaller impact in decision-making. On this reading, the PR phenomenon is explained by the overweighing of the payoffs in the pricing task insofar as both the response and the payoffs are expressed in monetary units. To give a final example, Daniel Kahneman’s and Amos Tversky’s (1979) *prospect theory* explains preference reversals on the basis of two phases affecting decision-making. In the editing phase individuals organise and reformulate the alternatives with the aim of simplifying the second phase of evaluation and choice. Thus, the decision taken is dependent on both the editing of the problem and the decision maker’s preferences.¹⁷⁵

Because they were trained in a different culture, the psychologists had no resistance in producing an interpretation of the PR phenomenon that was in sharp disagreement with a basic presupposition of economics, that individual preferences are pre-defined and fixed. Here it is how Tversky and Thaler put it:

First, people do not possess a set of pre-defined preferences for every contingency. Rather preferences are constructed in the process of making a choice or judgment. Second, the context and procedures involved in making choices or judgments influence the preferences that are implied by the elicited responses. In practical terms, this implies that behavior is likely to vary across situations that economists consider identical. (1990, p. 210)

Slovic considers that the plasticity of preferences includes “judgements and choices among options that are important, complex, and perhaps unfamiliar, such as gambles, jobs, careers, homes, automobiles, surgical treatment, and environments”. He adds “[i]n these decisions, preferences are not simply read off some master list but are constructed on the spot by an adaptive decision maker” (1995, p. 369). And this construction involves a manifold of mental operations:

Construction strategies include anchoring and adjustment, relying on the prominent dimension, eliminating common elements, discarding nonessential differences, adding new attributes into the problem frame in order to bolster one alternative, or otherwise restructuring the decision problem to create dominance and thus reduce conflict and indecision. As a result of these mental gymnastics, decision making is a highly contingent form of information processing, sensitive to task complexity, time pressure, response mode, framing, reference points, and numerous other contextual factors. (1995, p. 369)

9.7 DISCRIMINATING AMONG DIFFERENT HYPOTHESES

The plasticity of the experimental systems produced different and conflicting interpretations of the PR phenomenon. The PR phenomenon could be either explained as a violation of transitivity or as a violation of procedure invariance. However, the range of possible interpretations of the phenomena produced by plastic experimental systems can be narrowed down by the conduction of more experiments. This was the case with the PR research program. The co-existence of these interpretations triggered a new phase of experimental work intended to discriminate between them (Loomes *et al.* 1989, 1991; Schkade and Johnson 1989, Tversky *et al.* 1990).

Tversky *et al.* (1990) conducted an experiment to find out whether the PR phenomenon is caused by non-transitivity or by non-invariance. To this end, they extended the standard experimental design to include a certain amount of money (X). Under the new design the subject is first asked to state the lowest selling price for the P and the \$ bets. Subsequently, the decision maker is faced with pairs of bets and asked to select the preferred bet (\$, P or X). At the end of the session, one of these pairs is selected at random to determine the subjects’ payoff. If the selected pair belongs to the choice task subjects play the bet of their choice, if it belongs to the

¹⁷⁵ A more complete account of psychology explanations of PR can be found in Tversky and Thaler (1990) and Slovic (1995).

pricing task they play the bet for which they stated the higher selling price. Because the BDM procedure was not used, the hypotheses of non-independence and reduction were eliminated at the outset.¹⁷⁶ Were a PR to be verified it could only be caused either by intransitivity or a violation of procedure invariance.

Given that a PR occurs when the P bet is preferred to the \$ bet and the \$ bet is priced higher than the P bet ($P \succ \$$ and $P\$ > Pp$), intransitivity is the cause of the reversals if subjects prefer the \$ bet to the certain amount and the certain amount to the P bet (if $\$ \succ X$ and $X \succ P$, transitivity would imply that $\$ > P$ which is not observed when the PR occurs).¹⁷⁷ If procedure invariance is violated then there must be a discrepancy between choosing and pricing, which would result in the overpricing of the \$ bet and in the underpricing of the P bet. If the former observes, subjects prefer the certain amount to both the \$ bet and the P bet (that is $X \succ \$$ and $X \succ P$, and $P\$ > X > Pp$, which means that subjects price the \$ bet higher than the certain amount, but when facing a choice, they choose the certain amount). Procedure invariance would also occur in case of underpricing of the P bet, that is, when the P bet is preferred to the certain amount ($P \succ X$) but the price placed for the P bet is lower than the certain amount (because $P\$ > X > Pp$).¹⁷⁸ In whichever case, choice-based preference and value-based preference are different.

Having reproduced the standard reversal phenomenon (between 40% and 50%) that could not be attributed to an artefact caused by the elicitation procedure, they concluded that reversals were mainly caused by a failure of procedure invariance. They estimated that non-invariance explained about ninety percent of the reversals whereas the remaining ten percent of the preference patterns were explained by intransitivity (1990, p. 210). These results were also supported by a different experiment by Bostic *et al.* (1990).

In the economics camp, Loomes *et al.* (1989) designed an experiment to discriminate between the non-invariance explanation of preference reversals as put forward by Slovic and Lichtenstein (1983) and regret theory as developed by Loomes and Sugden (1983). They designed a special 'choice-only' experiment in which the subjects were confronted with three pairwise choices {P, \$}, {\$, C} and {C, P}, consisting of the lotteries P and \$, and of a sum of money with certainty, C. Under this design, regret theory predicts the occurrence of one preference reversal cycle: either the 'predicted cycle' $P \succ \$$, $\$ \succ C$ and $C \succ P$ or the opposite 'unpredicted cycle' $\$ \succ P$, $P \succ C$ and $C \succ \$$. According to the hypothesis of procedure non-invariance, there would be no reason to expect a cycle reversal because this design consists of a single task. To control the effect of the change in the design on the frequency of the phenomenon, the experimenters ran a standard

experiment and then imputed the absent choices, {C, P} and {\$, C}, in order to obtain the cycles of pairwise choice that would be obtained for those bets.¹⁷⁹

If reversals were to be caused by non-transitive preferences then there should be no difference between the two experiments. If however there were a procedural effect there would be a higher rate of cycles in the imputed choices of the standard experiment than in the actual choices of the choice-only experiment. Because the predicted pattern of reversals significantly outweighed the unpredicted pattern and there was no systematic difference between responses in the choice and the pricing tasks, the experimenters concluded that intransitivity is the primary cause of the phenomenon. However, the lower rate of reversals in the choice-only experiment was interpreted as an indication that the violation of procedure invariance could not be excluded as a cause of reversals. Loomes *et al.* concluded: "this is not to say that information-processing effects have no part to play in explaining choice under uncertainty; but it seems that these effects may not be the primary cause of preference reversals" (1989, p. 150). Loomes (1991) further noted that the differences in pricing between the \$ bet and the P bet in the standard experiment, which were more pronounced as the payoff of the \$ bet increased relative to the payoff of the P bet, could not be explained by non-transitivity. However, the primary cause of the phenomenon of preference reversals continued to be non-transitive preferences (Loomes *et al.* 1991).

The psychologists' explanation and the economists' explanation of PRs seem to be contradictory. While the psychologists consider the PR phenomenon to be mainly caused by the experimental procedure, economists consider the PR to be mainly caused by intransitive preferences. But both communities of researchers came to accept that non-invariance and intransitivity are causes of PR. What they disagree about is their relative impact on the phenomenon. As Loomes put it:

The overall conclusion appears to be that, although explanation 2 [violation of other axioms of EUT] can be rejected, failures of transitivity and failures of invariance both seem to be in evidence. These are not mutually exclusive, and their relative contribution to the preference reversal phenomenon has not yet been established beyond doubt (and, indeed, may never be precisely quantified, since the relative contributions may themselves vary with factors such as the parameters of the lotteries or the particular format in which the problems are presented). (1991, p. 602)

Thus, what seemed to be incompatible results are in fact complementary and reinforcing results. Intransitivity and the failure of procedure invariance are the two main sources of preference reversals whose relative contribution depends on the experimental procedure used to elicit preferences (cf. Starmer and Sugden 1998; Loomes 1998). This, in turn, provided further evidence for the high degree of sensitivity of experimental results on the experimental systems that produce them, and

¹⁷⁹ The choice experiment rendered more difficult the manifestation of the PRs. Whereas in the standard experiment a PR occurs whenever the valuations of the two bets are ranked differently, in the choice-only experiment a PR to occur requires that the valuation of the \$ bet is greater than the certainty amount C, and that the valuation of the P bet is lower than C.

¹⁷⁶ The argument is the same of Cox and Epstein (1989). The BDM mechanism could be eradicated because to "demonstrate preference reversals . . . it is not necessary to elicit the actual selling prices; it is sufficient to establish their order" (Tversky *et al.* 1990, p. 207, emphasis in original).

¹⁷⁷ The symbol \succ represents a relation of preference between bets and the symbol $>$ the ordering of prices.

¹⁷⁸ In this experiment subjects were paid for their participation but their payoffs were not contingent on subjects' decisions. Subjects were told to imagine that one pair of bets would be selected at random and that they would then have the chance to play either the bet they had chosen or the bet for which they had stated a higher cash equivalent. In one group, subjects knew that 15% of them would play the bet for real. No effect of monetary incentives was observed, however.

the need of bearing this in mind when interpreting them. Or to put it in another way, the influence of the larger socioeconomic context, laboratory and otherwise, on human behaviour renders the analysis of human behaviour a locally situated affair.

9.8 THE SIGNIFICANCE OF PR PHENOMENON TO ECONOMICS

The recognition that the violation of procedure invariance was a cause of the PR phenomenon had serious implications to economics. Rather than conceiving the individual as an agent who processes the various decision tasks on the basis of stable and well-defined preferences and universal principles, economists started to consider that the decisions individuals take are influenced by the decision-problems they face which determine both their preferences and the chosen decision heuristics.

By focusing attention on particular axioms such as independence and transitivity, we have overlooked an even more fundamental assumption, which most economists seem to take for granted, but which is almost certainly false: namely, that people come to problems armed with a clear and reasonably complete set of preferences, and process all decision tasks according to this given preference structure. But I believe that the reality is very different, and that most people's preferences are generally imprecise and in many respects incomplete, with the result that they are liable to process different decision tasks in rather different ways. (Loomes 1999b, F37)

The possibility that decision-rules may be contingent to decision-problems in turn suggested a novel research problem: the study of the processes of decision-making. Loomes put the implication of the PR research program to economics in the following terms:

Instead of continuing to try to devise some general theory of an essentially conventional (e.g. axiomatically based) form, perhaps we should switch our attention and our efforts to understanding more about the processes by which people select and apply rules/strategies for dealing with particular forms of decision problem. As part of *this* agenda, we would need to examine how robust such rules/strategies are, how they may evolve or be modified in response to feedback or experience, and in what ways the predictions (and, where relevant, the prescriptions) that follow from them may diverge from those derived from more conventional models. (1998, p. 486 emphasis in original)

The implications that followed to experimental economics were, instead:

The challenge is not to refashion existing experiments to incorporate even tighter controls until they succeed in generating results consistent with conventional theory by reducing participants to barely more than 'zero intelligence traders'. Rather, the real challenge for the future is to devise experiments which allow for the heterogeneity of human behaviour, and to develop techniques which give greater insights into the interactions between people's imprecise basic values and the environments in which they have to operate, tracing how they construct their responses and/or modify them in the light of experience. (1999b, F44)

It is clear that Loomes is here advocating the conduction of 'behavioural experiments' to study of the processes of decision-making. A high level of 'agency' on the part of the experimental subjects must be allowed in economic experiments in order to examine the interaction between individuals' basic values and the contexts wherein they operate. 'Technological experiments' are not suitable to this end. Instead, they provide knowledge about the ways by which individual preferences and behaviour can be induced and controlled by highly structured contexts of interaction.

The profession at large still neglects the implications of experimental evidence to EUT. It is still the case that EUT remains as the main theoretical framework of decision-making under risk. The partial non-expected utility theories that economists developed to accommodate the anomalies to EUT still presuppose that preferences are fixed and that agents maximise preferences.¹⁸⁰ However, the reluctance of economists in accepting experimental evidence cannot be attributed to the experimental method of economics. It is instead a resistance towards rejecting a well-established theory of economics.

9.9 THE RE-EMERGENCE OF THE METHODOLOGICAL DEBATE

A recent and radical response on the part of the experimental economists denies the significance of the PR phenomenon to economics arguing that the precepts of experimental economics are not satisfied in the PR experiments. The claim is that economic theory applies to contexts where these precepts are verified and that the PR phenomenon is not observed therein. But this strategy implies a far more delimited scope of economics which is anyway refuted by the very practice of the economists.

Glenn W. Harrison (1994) conducted an experiment to demonstrate that the PR phenomenon does not occur when the conditions for a valid controlled experiment in economics are satisfied (see Ch. 5). According to Harrison, these conditions were not met in previous experiments because dominance was violated. The claim is that the incentive structure did not satisfactorily compensate for the subjective costs associated with the experimental tasks. Even though the criticism targeted all PR experiments, Harrison focused on the experiments by Grether and Plott (1979), Reilly (1982) and Pommerehne *et al.* (1982) to demonstrate that in these experiments the opportunity cost of an inconsistent report, i.e. the expected income that the individual is prevented from receiving by reporting inconsistently, was too low. On this reading, preference reversals occurred because the opportunity cost of mistakes in the valuation task was too low.¹⁸¹ The underlying presupposition is that if subjects had the opportunity to experience the consequences of their mistakes the inconsistency could have been corrected and the anomaly would have been eliminated. Harrison designed an

¹⁸⁰ A very comprehensive account of recent developments in non-expected utility theories can be found in (Starmer 2000).

¹⁸¹ The presumption is that subjects report true preferences in the choice task and that the reversals occur when subjects place higher prices for the \$ bet. The opportunity cost of false reports is then measured on the basis of the valuation elicited for the \$ bet, the mistakenly highly priced bet.

experiment in which the opportunity cost of errors was increased by changing the scale of reporting selling prices (these had to be stated in increments of 25, 50, 100 or 200 cents) and by widening the interval of the expected values of the lottery pairs. As expected, the rate of reversals declined (e.g. a rate of 10% was observed when the disparity was of 100 cents). In this way the results provided evidence that inconsistent behaviour can be corrected if individuals suffer the consequences of inconsistency. However, the changes in the design rendered the manifestation of the PRs less likely. The reason is that the change in the scale of the selling prices may have hidden reversals that otherwise would have been revealed. Moreover, it could be argued that the change in the incentive structure was not salient to the subjects because they had no feedback on the consequences of their actions (c.f. Tammi 1999a, p. 364-5 and references therein). In sum, since the design of the experiment was substantially different from standard experiments the decline in the reversals cannot be directly imputed to the increase of the opportunity costs. Moreover, recent studies on the effect of financial incentives have supported the view that the simple inclusion of substantial financial incentives does not eradicate violations to EUT (Beattie and Loomes 1997, Camerer and Hogarth 1999).

Peter Bohm (1994) showed that in the context of markets for claims of non-negligible magnitudes, competent decision-makers (i.e. third-year business school students specialising in finance) do not incur as often in PR (only 15%).¹⁸² Along the same line of reasoning, Cox and Grether (1996) devised an experiment which demonstrated that the PR phenomenon could be eradicated in a market with high monetary incentives, immediate feedback and repetition. They found that market institutions with strong incentives and immediate feedback are not sufficient to make reversals disappear. The reduction of reversals requires that subjects must have also the opportunity to learn with experience. Only in repetitive markets does the rate of reversals reduce and does the asymmetric pattern disappear.

[W]e are inclined to rule out feedback and incentives alone as causes of discrepancies between the results of market and individual choice experiments. With repetition, the rate of preference reversals fell substantially with all methods and the asymmetry between rates of predicted and unpredicted reversals generally disappeared. This suggests that the repetitive nature of the tasks in market experiments in conjunction with feedback is an important factor. (op. cit., p. 402)¹⁸³

In sum, what seems to be crucial for the elimination of the reversals is that subjects are given the opportunity to learn the incentive structure and to act in conformity.

¹⁸² This experiment involved decisions between claims at different future dates. The PR would occur in the presence of an inconsistency between a time preference revealed by choice and pricing of claims of amounts payable at different future dates (e.g. choice between \$200 from three months now and \$240 from fifteen months now).

¹⁸³ This had already been ascertained by previous market experiments which failed to eliminate PRs (e.g. Knez and Smith 1987).

A similar result was obtained by Yun-Peng Chu and Ruey-Ling Chu (1990) in a market experiment where an arbitrager could take advantage of subjects' inconsistent behaviour. In this experiment, subjects were confronted with the same pair of bets about which they had to indicate a preference and to state monetary values for each one of them. Arbitrage would take place whenever the subject incurred in a preference reversal. In the case of the predicted reversal, the experimenter would sell the \$ bet at the reservation price $P_{\$}$, exchange the P bet for the \$ bet, and then buy the P bet at the stated price P_P . Because $P_{\$} > P_P$, arbitrage would money-pump the subjects.

It is, therefore, not surprising that the reversals were eliminated in a market where subjects were exposed to repeated transactions that cause them monetary losses. In short, arbitrage was very successful in making experimental subjects to conform to the predicted outcome. But arbitrage alone did not suffice to reduce PR. The results of Chu and Chu experiment suggested that the main factor responsible for the eradication of the PR is the continued exposure to arbitrage during a number of decision periods for the same pair of bets. This is, in fact, supported by the arbitrage experiment conducted by Berg *et al.* (1985) that failed reducing the frequency of reversals in an experiment where subjects were exposed to arbitrage only once (though it did succeed in reducing the size of error).

These experiments were presented as evidence to the effect that in economic contexts experimental subjects behave as predicted by economic theory. However, the experimenters fail to understand that the behaviour observed was substantially induced by experimental design. In other words, because the agency of the experimental subjects was significantly reduced, these experiments do not provide meaningful information about human preferences and behaviour. Under such circumstances, subjects could not have acted otherwise. Moreover, the difficulty in obtaining such behaviour in the laboratory show that the behaviour elicited observes only in very structured environments.

9.10 UNIFYING THE EXTANT RESULTS INTO A COHERENT ACCOUNT

That preference reversals do not occur in some experimental markets while they occur in a variety of other experimental situations is not a controversial observation. But it has been used in support of conflicting conclusions. As it was shown above, psychologists and an increasing number of experimental economists came to see the evidence produced by the PR research program as refuting evidence to EUT and, most damaging, a challenge to a basic presupposition of economics, the stability of individual preferences. The recent eradication of PR in some market experiments, on the contrary, inspired a new interpretation of the experimental evidence. A unifying reading of the experimental results by Charles Plott (1995) not only attempted to save EUT but it also attempted to save the basic presupposition of economics. This was accomplished by the 'discovered preference hypothesis' (DPH).

The DPH is an intermediate mode of inference that takes experimental data as evidence of a more general pattern of human behaviour that extends beyond the

behaviour observed in the experimental world. It identifies a behavioural pattern that applies generally rather than to well defined and concrete situations in the real world. The main claim of the discovery preference hypothesis is that rather than possessing inconsistent and unstable preferences, individuals have consistent and stable preferences which are not always revealed in the decisions they make. Preferences are instead ‘discovered’ as individuals come to know what they want through practice:

The hypothesis suggests that attitude like expectations, beliefs, risk-aversion and the like, are discovered, as are other elements of the environment. People acquire an understanding of what they want through a process of reflection and practice. In a sense, they do not know what they want and it may be costly, or even unpleasant, to go through the process of discovery. Attitude discovery is a process of evolution which has a direction, and in the final stage results in the ‘discovery’ of a consistent and stable preference. (op. cit., p. 227)

The DPH takes experimental data as evidence pertaining to situations that deviate from those to which rational choice theory applies. Rational choice theory does not apply to situations where subjects are inexperienced with the task at hand or unable to take into account the behaviour of others in the decision-making process. However, the deviations to rational choice theory come in different degrees. The greatest deviation occurs when there is an absolute lack of experience and a very limited awareness about the immediate environment or the consequences of the decisions to be taken. The individual is nevertheless purposeful and optimising. He is simply unable to avoid some inconsistency when dealing with new tasks. Under these circumstances, “responses are ‘instantaneous’ or ‘impulsive’, reflecting whatever may have been perceived as in self-interest at the instant” (Plott, 1995, p. 226). Nonetheless, “systematic aspects of choices might exist, reflecting attention and perception, but they might not make sense when viewed from the perspective of a preference based model” (op. cit., p. 227). At an intermediate level of rationality, the individual has some experience so that “choices begin to reflect and incorporate an awareness of the environment, and can be recognised by an ‘outsider’ as a stable form of ‘strategy’ or ‘decision’” (ibid). But it is only when choices comprehend expectations about the behaviour of others that individual behaviour may be regarded as rational. The upshot is that “under conditions of substantial incentives, and with the accumulating information that is obtained from the process of choice, the attitudes stabilize in the sense of a consistent decision rule, reflecting the preferences that were ‘discovered’ through the process” (op. cit., p. 228).

The DPH can save rational choice theories but it does not come without a cost. DPH suggests that rational choice theory is of a very limited applicability. They apply to situations where the individual is experienced and capable of anticipating the behaviour of others which, as Plott acknowledges, pertain to very structured environments. The analysis of experimental economics supports this conclusion. It shows that rational behaviour, i.e. the tendency to act in conformity with the experiments’ incentive structure, occurs in contexts where the monetary reward is the main motive and subjects can experience the financial consequences of their actions and thereby have the opportunity to learn and to adjust behaviour accordingly. To put it in another way, rational behaviour occurs in experimental contexts in which ‘human

agency’ is low because there is a single salient course of action. In the standard PR experiment the level of ‘human agency’ was high because there was not a salient course of action. Subjects had simply to choose and price lotteries. But when the subjects suffered monetary losses by making inconsistent decisions they correct them. But it is clear by now that the fact that the experimental design succeeded in triggering consistent behaviour does not mean that it also succeeded in eliciting subjects’ true preferences. It simply means that consistent preferences were successfully induced in the laboratory.

At any rate, the DPH does raise two important and interrelated debates, one pertaining to the scope of economics and the other to the method of experimental economics. Even though time and length constraints prevent a more in-depth contribution to these debates, I can nonetheless hint at the response afforded by SEEE.

Plott seems to suggest that economic theory purportedly applies to structured environments where subjects possess sufficient experience to make rational choices. In this regard, an obvious remark is patent. The very practice of economic, theoretical and empirical, denies this. Indeed, the most progressive developments of economics in recent years have been observed in areas of research that deviates from the domain of rational choice theory (e.g. behavioural economics). This not to mention that there are phenomena of economic interest that do not involve experience, such as decisions that are rare and/or irreversible (e.g. childbearing, marriage, job-taking, decisions relating to health, education, and so forth).

The restricted view on the scope of economics is often associated with a restricted view on experimentation. The most explicit of which is eloquently defended by Ken Binmore:

My own experimental papers therefore insist that economic theory should only be expected to predict in the laboratory if the following three criteria are satisfied:

The problem the subjects face is not only ‘reasonably’ simple in itself, but is framed so it seems simple to the subjects;

The incentives provided are ‘adequate’;

The time allowed for trial-and-error adjustment is ‘sufficient’. (1999, p. F17)

But as explained before, the simplicity or transparency of the experimental task and the imposition of monetary incentives are requirements that are meant to achieve control over the economic experimental systems (see Ch. 6). They ensure that subjects understand the experimental situation in the intended way and act accordingly. The inclusion of repetition, on the other hand, depends on the problem-situation the experiment intends to solve for there are phenomena of economic interest that do not involve experience. That is, the requirement of experimental control does not demand time for trial and error.¹⁸⁴ To put it in another way, the precepts of experimental

¹⁸⁴ A similar argument may be found in Cubitt et al. (2001). They argue that single-task individual-choice designs may be effective in exercising control over experimental systems and therefore they may do without market institutions, feedback mechanisms and high incentives. The argument is that these are transparent and simple experiments which may economise on time required for learning and rewards for mental effort on deliberation.

economics are intended as procedures that aim at achieving control over the economic experimental systems rather than boundary conditions of experimental economics. Simplicity is a requirement that any experiment must fulfil for it raises the intelligibility of the experimental results. The use of monetary incentives is a specific requirement of experimental economics insofar as it is instrumental to induce in the laboratory interesting behaviour to economics.

However, as it should be clear by now, the fact that experimental subjects behave in accordance with the incentive structure of economic experiments is not very informative. The use of monetary incentives in economic experiments performs other epistemological goals. It allows learning how to control human behaviour for intended purposes and to identify the factors that cause deviations to income maximising behaviour. The accumulation of laboratory anomalies that do not easily go away with high incentives and the work required so as to induce the 'right' behaviour in the laboratory advice caution in what regards the scope of both economics and experimental economics. Human behaviour might be very different than that predicted in contexts, experimental and otherwise, where economic theory supposedly applies.

Plott's discovered preference hypothesis is meant as an alternative to the opposite hypothesis put forward by the psychologists that claimed that preferences are constructed. The low level of 'human agency' of the experiments that corroborate it does not afford a strong empirical support. The accumulation of anomalies by experimental economics and psychology instead constitute substantial evidence for the fuzziness and context-dependence of preferences.

9.11 CONCLUSION

The research program of preference reversals is very illuminative in illustrating how the social dimension of knowledge production and the methodological specificities of experimental economics interact.

The research program produced socially robust results because it took place in a context that favoured effective criticism (see Ch. 3). These results were the outcome of a prolific dialogue between scientists educated in different cultures. The experimenters were in a position to pose serious challenges to the rival standpoints and compelled to generate strong responses to the challenges targeted at them. Moreover, this process of social criticism involved a wide community of researchers that included practitioners with expertise in different and disparate resources of practice, both theoretical and experimental. As a result, it is now well-established what are the factors preference reversals are vulnerable to. More importantly, the dialogue between scientists from different cultures strongly contributed to the identification and scrutiny of a major presupposition of economics. In fact, it is now part of the subject matter of the field of behavioural economics. Behavioural economics, originally inspired by the accumulation of experimental anomalies to economic theory, is now concerned with the inclusion of the experimental underpinnings in the models of human behaviour (Camerer and Loewenstein 2004).

But the importance of the social dimension of knowledge production is not circumscribed to the definition of interesting problem-situations and to the generation of robust problem-solutions. The consideration of the social and the individual disposition to accept results consistent with previous expectations and to reject results that conflict with them conveys a higher epistemic value to the results that challenge established knowledge.¹⁸⁵ These possess an added epistemic value insofar as they manifest a strong resistance to substantial efforts to eliminate them. Thus, the fact that some practitioners in the field are now engaged in a new line of research is therefore a very significant contribution of the PR research program.

The PR dispute has been interpreted as a demonstration of dogmatism on the part of the economists. Experimental economists were accused of devoting too much time attempting at eradicating a phenomenon that was well established by the psychologists. The theoretical economists were accused of ignoring experimental anomalies to economic theory. The analysis developed here shows that the 'dogmatic' attitude of the experimental economists has epistemic value. Not only was it important to improve the understanding of the anomalies, but it was also crucial for accepting them. The persistent and failed attempts at materially eliminating or explaining the phenomenon away ultimately convinced the economists of its significance. This transformation of beliefs could never be operated by evidence produced by psychologists. However, the epistemic value of 'dogmatism' requires an effective critical community in which scientists are compelled to pose and answer critical challenges.

The study of the PR research program allows deriving more general conclusions about laboratory experimentation in the human sciences. First, the behaviour observed in the laboratories of the human sciences is very contingent on the experimental design. Second, in the human sciences clear-cut results can only be produced by rigid experimental systems. The use of plastic experimental system requires a sequence of experiments each of which improves the understanding of the previous results. Third, the virtual possibility of obtaining any result means that the analysis of the problem-solutions arrived at must take into account the participation of the human subjects in their production. Fourth, the high level of vulnerability of the experimental results to scientists' priors reinforces the importance of the collective process of criticism which is facilitated in the human sciences. The process

¹⁸⁵ Indeed, this has been the object of analysis of psychology. The research carried out has shown that once in possession of strong beliefs, individuals do not pay enough attention to, or may even misinterpret, new information contradicting these beliefs. This tendency is sharper in the presence of complex and ambiguous evidence. A review on belief perseverance and confirmatory bias can be found in Rabin (1998). In the same review a direct reference is made to scientists' biases and, in particular, to economists' reluctance in accepting behavioural evidence. In this regard it is worth quoting the following citation:

With confirming evidence, we suspect that both lay and professional scientists rapidly reduce the complexity of the information and remember only a few well-chosen supportive impressions. With disconfirming evidence, they continue to reflect upon any information that suggests less damaging "alternative interpretations". Indeed, they may even come to regard the ambiguities and conceptual flaws in the data opposing their hypotheses as somehow suggestive of the fundamental correctness of those hypotheses (Lord et al. 1979 in Rabin 1998, p. 28, emphasis omitted).

and products of experimental activity are accessible to a wide community of researchers as well as to the public in general.

The normative implication of the present chapter is that the social organisation of science should encourage dialogue between scientists belonging to different research traditions. More stringent and robust results to interesting problem-situations can be expected from processes of knowledge production that articulate items from disparate cultures. The progress of science ultimately depends on the identification and exploration of unconsciously held beliefs. To this end, scientific practice would be more fruitfully organised around problem-situations, the resolution of which would have to be effective in mobilising the relevant resources irrespective of their origin. To put it in another way, SEEE strongly discourages problem-solving restricted by the borderlines of well-delimited fields of research. This is particularly pertinent in the human sciences given the multidimensionality and complexity of human behaviour. This is what has been demonstrated by the foregoing analysis of experimental economics.

CHAPTER 10

CONCLUSION: WHAT ABOUT THE SOCIAL EPISTEMOLOGY OF EXPERIMENTAL ECONOMICS?

This thesis analysed the experimental process of knowledge production. It investigated how scientists build their confidence in knowledge generated by a process in which both the means and the outcomes of knowledge production are re-constructed.

The study of experimental practice in the natural and human sciences supports the view that scientists are convinced that they have produced the phenomenon of interest when they achieve a *three-way coherence* between the three components of the experimental system: the experimental procedure, the instrumental model and the phenomenal model. When the three-way coherence is achieved, experimenters believe that they have created an experimental system that succeeded in producing the phenomenon of interest. The relation of coherence among the three components of the experimental system justifies belief in the experimental results because the three-way alignment supports each one of them and thus the experimental result conveyed by the phenomenal model. This was the underlying principle of the *argument from coherence* that justifies the way by which experimenters form belief in experimental results. However, it was also noted that the three-way alignment is not sufficient to justify belief in experimentally generated knowledge. Two additional arguments were presented that reinforced the epistemic value of the three-way coherence.

The *argument from materiality* asserts that the direct engagement of the subject matter in knowledge production (both in the natural and human domains) renders experimental results and the coherences supporting them non-trivial achievements. The coherent problem-solutions arrived at carry knowledge about the subject under scrutiny because scientists cannot fully control it to meet their prior expectations. However, the argument from materiality does not satisfactorily account for experimenters' confidence in experimental results. The participation of the subject matter might still be severely constrained by the problem-situation at hand or by the plasticity of the experimental systems. The *argument from sociality* asserts that the social dimension of knowledge production encourages the generation of fruitful problem-situations and reliable problem-solutions by bringing to the production process a vast number of resources of practice. The three arguments in conjunction lead to a broader conclusion: the greater the number and the greater the heterogeneity of the resources (material, conceptual and social) involved in knowledge production, the higher the epistemic status of the relations of coherence established given that they are the result of practices that have explored relevant courses of action to the resolution of interesting problem-situations.

The research carried out here, however, is not immune to the problem it identified. The study of experimental economics was based on an epistemological framework purposefully designed for the analysis of the processes and products of experiments – the Social Epistemology of Experiment (SEE). The strategy used to tackle this problem was the same as that of the experimental scientists. SEE is based on robust three-way coherences that support the account of the processes in which scientists justify their beliefs in knowledge generated by experimental means. To put it another way, the research carried out conforms to both the descriptions and the prescriptions of SEE. The demonstration of this claim is thus the natural topic of this concluding chapter.

10.1 APPRAISING SEE

The study of experimental economics had recourse to an analytical framework that was purposively constructed to that end. However, it can be shown that SEE is a coherent resolution that circumvents the epistemic problems that arise when both the means and the outcomes of knowledge production are re-constructed. To show this, enumerated below is the list of the coherence procedures followed and the results achieved. The conformity to SEE's prescriptions may be independently appraised.

- i) *SEE is a robust and coherent analytical device.* The construction of SEE relied upon various and heterogeneous items of scientific culture of philosophy, history, sociology, economics, and psychology, all of which concur with the account of scientific experimentation put forward.
- ii) The knowledge production process that generated SEE involved the participation of the object of scrutiny. In other words, SEE conveys the 'material agency' of the subject matter, i.e. actual experimental practice.
- iii) SEE is *stringent* framework of analysis insofar as it has generated a wide range of clear-cut results.
- iv) *SEE has a high level of technological applicability* insofar as it is applicable, as was applied, to various episodes of experimental practice, both in the natural and in the human sciences, as well as to various processes, methods and results of experiments.
- v) *SEE has a potential high level of social robustness* insofar as it is amenable to the scrutiny of a wide and heterogeneous audience. However, this test is still to be carried out.
- vi) The major contribution of SEE is, however, the identification of new areas of research in experimental economics, in economics in general and in the studies of science. Overcoming the limitations of the present conditions of scientific practice is after all the central concern that pervades SEE.

10.2 APPLYING SEE TO EXPERIMENTAL ECONOMICS

The application of SEE to experimental economics turned out to be very fruitful. A broad portrait of experimental economics was provided along with a more fine-

grained analysis of particular episodes and results. These included the scrutiny of arguments intended to justify the use of the experimental method in economics, research programs, experiments, experimental procedures, and social practices. Not only did SEE bring novel insights into ongoing methodological debates and controversies, but it also highlighted novel problems that deserve the attention of students and practitioners of experimental economics. The successful implementation of SEE thereby provides support for the adequacy of the framework for both the descriptive and normative analysis of the processes and products of experiments. However, the application of SEE to experimental economics identified features common to the scientific enterprise, features specific to the experimental process of knowledge production, and features characteristic of the experimental field of economics. Each of which are described below.

SEE ON SCIENTIFIC PRACTICE

- i) The establishment of a new field of research is the outcome of a long process of knowledge production that involves a wide community of researchers and multiple items of various scientific cultures.
- ii) In contrast, work carried out in well-established fields of research is more efficient and effective because it does not have to be self-justificatory.
- iii) However, the new fields of research carry a high potential for the growth of knowledge. The work done in well-established fields consists first and foremost in furthering the articulation of already fairly established items of scientific culture.
- iv) The scepticism new proposals encounter has epistemic value. It motivates scientists to produce sound and robust arguments that can face up the criticisms of the sceptics and ultimately win collective acceptance.
- v) Interdisciplinary work (or better the disregard for disciplinary boundaries) is fundamental to identifying unconsciously held beliefs and thereby novel areas of research.
- vi) Scientific institutions must actively encourage the removal of disciplinary boundaries impeding growth of knowledge because the social, psychological and the cognitive characteristics of communities of scientists do not necessarily promote it.

SEE ON THE EXPERIMENTAL PROCESS OF KNOWLEDGE PRODUCTION

- vii) Control is the hallmark of scientific experimentation and it is the main factor responsible for the local and situated nature of experimental knowledge.
- viii) The central epistemic question of scientific experimentation concerns the plasticity of experimental systems.
- ix) The plasticity of experimental systems requires the appraisal of the extent to which experimental results are the outcome of scientists' material and conceptual manipulation rather than the agency of the subject matter, i.e. the material world in the natural sciences and the human world in the human sciences.

- x) The social context in which knowledge production takes place, i.e. the social world, can substantially improve the rigidity of experimental systems. This can be done by a series of experiments each of which improves upon the previous results.
- xi) There is no dividing line between the natural and human sciences as far as the epistemic value of experiments is concerned. The direct participation of the subject matter in the experimental process of knowledge production gives experiments epistemic superiority as compared to other methods of knowledge production.

SEE ON EXPERIMENTAL ECONOMICS

- xii) In economics, control is exercised via a monetary reward structure that prizes economical successful decisions made in a context of relative anonymity, and via institutional rules that govern the laboratory contexts of interaction.
- xiii) The high level of control exercised in economic experiments reduces the participation of experimental subjects and thus the level of human agency in the experiment.
- xiv) The control exercised in economic experiments, however, renders economic experiments fairly transparent experimental systems from which inferences can be made with a high degree of confidence.
- xv) The level of control exercised in laboratory experiments can be lessened and the resulting reduced level of transparency improved by new experiments.
- xvi) The fact that experimental results are very open to the scrutiny from a wide audience contributes to the high level of reliability of the experimental process of knowledge production and to the participation of other scientists.
- xvii) The applicability of the results of economic experiments to concrete situations in the real world derives from the possibility of creating (in both the laboratory and the non-laboratory worlds) social contexts that promote socially-desirable goals.
- xviii) Economic experiments generate knowledge about the factors that affect the manifestation of self-interested income maximising behaviour and the means by which human behaviour may be controlled for socially-desired goals.

10.3 THE NOVEL CONTRIBUTIONS OF SEE

The major contributions of SEE are, however, the novel insights it has brought to the study of scientific experimentation in economics and the new areas for research it identified.

THE MATERIALITY AND THE SOCIALITY OF SCIENTIFIC EXPERIMENTATION

SEE selects the notions of 'materiality' and 'sociality' as the main attributes around which scientific experimentation should be accounted for and appraised. The analysis of the 'materiality' of economic experiments pointed to the importance of assessing the contribution of the 'material' component of the experimental

systems, i.e. the participation of human subjects in economic problem-solving. Even though the participation of human subjects is the distinguishing feature of scientific experimentation in the human sciences, students of experimentation have not explicitly addressed the significance of this participation. In addition, the differentiated role this participation plays in economic experiments has not been explored. Experimental economists, in turn, do not seem to be conscious of the fact that the more control is exercised over human motivations and actions the less support the experiments provide to claims about them. It shows instead how the observed behavioural patterns can be generated in experimental worlds.

The social aspect of knowledge production turned out to be epistemically very relevant to experimental economics. The epistemic value of economic experiments is derived significantly from the high degree of susceptibility of experimental processes and products to public scrutiny. This susceptibility improves not only the reliability of the experimental process of knowledge production but it also improves the social robustness of the experimental results. The wide accessibility of experimental results to various publics carries a high potential for the generation of interesting problem-situations and robust problem-solutions. Indeed, it is this aspect of the experimental process of knowledge production that underlies the incremental way in which experimental knowledge is generated in economics. The collective nature of experimentation in economics, in addition, undermines the significance of the methodological and epistemological problems associated with the high degree of control.

The study of experimental economics undertaken thus far has revolved around the comparative analysis of experiments with other scientific products and on the relation between experiment and theory, or some target system the experiment allegedly refers to. SEE highlights instead the study of the sources of epistemic value of economic experiments: first, the actual participation of human subjects in experiments and the causal efficacy of institutions in constraining and enabling human behaviour; and second, the social dimension of knowledge production and its role in overcoming the methodological difficulties associated with experimenting with humans and the effect of individual and collective biases in knowledge production.

THE TECHNOLOGICAL AND THE BEHAVIOURAL EXPERIMENTS OF ECONOMICS

SEE identified the technological and behavioural experiments of economics, the epistemological functions of which have not been fully acknowledged by the experimenters themselves. Economists do not seem to be fully aware that some experiments in economics provide knowledge about how best to control human behaviour for specific purposes. Inferences about human preferences and behaviour can only be produced by experiments that provide experimental subjects with human agency. Indeed, both experimental economists and the students of experimentation have not yet fully recognised the epistemic role of this latter kind of experiments, i.e. the behavioural type of economic experiments.

Experimenters and students of experimentation have to date focused on the success stories of experimental economics, which concern the technological experiments of economics. These seem to accomplish the traditional aspirations of

the discipline of economics, i.e. to provide guidance for policy-making. In fact, economic experiments can do more than that. They can also be tools of social engineering insofar as they can be used for the design of socioeconomic institutions to be implemented in the real world. However, the behavioural experiments of economics cannot accomplish such goals. The results they produce do not have a direct application to concrete real world problem-situations. Nonetheless, they have generated important information about human behaviour which is also relevant for policy-making. The insights of behavioural experiments are now starting to feedback onto economics with as yet unknown outcomes. The high potential for generating knowledge that challenges the basic presuppositions of conventional economic theory suggests, however, that the behavioural experiments of economics can have a substantial impact on the overall conception of the discipline. SEE therefore strongly supports the methodological and the epistemological study of behavioural experiments and their impact on economics.

OTHER AREAS OF FRUITFUL RESEARCH...

SEE identifies other fruitful problem-situations. Particularly promising investigations concern the study of experimental work at the cross-section of various disciplines. The very emergence of the experimental method in economics derives from the interdisciplinary research carried out in the USA right after World War II that joined economists, psychologists, cognitive scientists, physicists, and mathematicians, among other specialists. At present, the recent collaboration between economists, psychologists and neuroscientists suggests the emergence of another exciting field of research with a high potential to feedback upon the discipline of economics.

SEE, in addition, highlights the relevance of the study of the organisation of institutions of science to further both the growth of knowledge and the values of good science. The increasing complexity of science and its growing influence on human and non-human affairs make this study a pertinent research question for all sciences. In economics, in particular, the inescapable and growing influence of the products of economic science on the wider social and political dimensions of human life (as well as the reverse, i.e. the influence of the wider context on the practice of the economists) makes this an ever-pressing issue. These issues fall within the realm of the sociology of economics which have already been, and are being, addressed by economists and sociologists. What SEE stresses is the importance of studying the mechanisms that can create more favourable conditions for multicultural dialogue and the creation of more effective critical communities. This dialogue and criticism are required not only for raising interesting but also socially responsible problem-situations that take account of those whose interests tend to be underrepresented in scientific practice – the less privileged and the future generations. This is, in the end, the concern that guides, or at least should guide, the practice of scientists and that of the students of science.

REFERENCES

- ACHINSTEIN, PETER 1968. *Concepts of Science: A Philosophical Analysis*. Baltimore: John Hopkins Press.
- ACHINSTEIN, PETER AND OWEN HANNAWAY (eds.) 1985. *Observation, Experiment and Hypothesis in Modern Physical Science*. Cambridge, Mass.: MIT Press.
- ACKERMANN, RICHARD 1989. "The New Experimentalism". *British Journal for the Philosophy of Science*, 40: 185-90.
- ALLAÏS, MAURICE 1953. Le Comportement de l'Homme Rationnel devant le Risqué: Critique des Postulats et Axioms de l'École Americane. *Econometrica*, 21: 503-46.
- ARROW, KENNETH J., SAMUEL KARLIN AND PATRICK SUPPES (eds.) 1960. *Mathematical Methods in The Social Sciences*. Stanford: Stanford University.
- ARROW, KENNETH J., ENRICO COLOMBATTO, MARK PERLMAN, AND CHRISTIAN SCHMIDT (eds.) 1995. *The Rational Foundations of Economic Behaviour*, IEA Conference, London: Macmillan Press.
- BACON, FRANCIS 1859 [1620]. *Novum Organum*. Chicago: Open Court.
- BANKS, JEFFREY S., JOHN O. LEDYARD AND DAVID P. PORTER 1989. "Allocating Uncertain and Unresponsive Resources: An Experimental Approach". *RAND Journal of Economics*, 20: 1-25.
- Bardsley, Nicholas 2005. "Experimental Economics and the Artificiality of Alteration". *Journal of Economic Methodology*, 12: 239-51.
- BARNES, BARRY 1974. *Scientific Knowledge and Sociological Theory*. London: Routledge & Kegan Paul.
- BARNES, BARRY 1977. *Interests and the Growth of Knowledge*. London: Routledge and Kegan Paul.
- BARNES, BARRY 1982. *T.S. Kuhn and Social Science*. London: Macmillan
- BATENS, DIDERIK AND JEAN PAUL VAN BENDEGEN (eds.) 1988. *Theory and Experiment: Recent Insights and New Perspectives on Their Relation*. Dordrecht: D. Reidl.
- BEATTIE, JANE AND GRAHAM LOOMES 1997. "The Impact of Incentives Upon Risky Choice Experiments". *Journal of Risk and Uncertainty*, 14: 155-168.
- BECKER, GORDON M., MORRIS H. DEGRÖOT, AND JACOB MARSCHAK 1964. "Measuring Utility by a Single Response Sequential Method". *Behavioral Science*, 9: 226-232.
- BELL, DAVID 1982. 'Regret in Decision Making under Uncertainty'. *Operations Research*, 20: 961-81.
- BERG, JOYCE E., JOHN W. DICKHAUT AND JOHN R. O'BRIEN 1985. "Preference Reversal and Arbitrage". In V. L. Smith (ed.), pp. 31-72.
- BERNOULLI, DANIEL 1954 [1738]. "Exposition of a New Theory on the Measurement of Risk". *Econometrica*, 22: 23-36. [transl. from Italian]
- BHASKAR, ROY 1975. *A Realist Theory of Science*. Brighton: Harvester.
- BIJKER, WIEBE AND THOMAS P. HUGHES AND TREVOR PINCH (eds.) 1987. *The Social Construction of Technical Systems: New Directions in the Sociology and History of Technology*. London: MIT Press.
- BINMORE, KENNETH 1994. *Playing Fair*. Cambridge, Mass.: MIT Press.
- BINMORE, KENNETH 1999. 'Why Experiment in Economics?' *The Economic Journal*, 109: F16-F24.

- BINMORE, KENNETH, AVNER SHAKED AND JOHN SUTTON. 1985. "Testing Noncooperative Bargaining Theory: A Preliminary Study". *The American Economic Review*, 75: 1178-80.
- BLACK, MAX 1962. *Models and Metaphors*. Ithaca, NY: Cornell University Press.
- BLAUG, MARK 1992 [1980]. *The Methodology of Economics: Or How Economists Explain*. 2nd edition. Cambridge: Cambridge University Press.
- BLOOR, DAVID 1976. *Knowledge and Social Imagery*. London: Routledge.
- BLOOR, DAVID 1983. *Wittgenstein: A Social Theory of Knowledge*. London: Macmillan.
- BLOOR, DAVID 1991. *Knowledge and Social Imagery* (2nd ed.). Chicago: University of Chicago Press.
- BLOUNT, SALLY 1995. "When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences." *Organizational Behavior and Human Decision Processes*, 63: 131-144.
- BOHM, PETER 1994. "Time Preference and Preference Reversal among Experienced Subjects: the Effects of Real Payments." *Economic Journal*, 104: 1370-1378.
- BOHNET, IRIS AND BRUNO S. FREY 1999. "Social Distance and Other-Regarding Behavior in Dictator Games: Comment." *The American Economic Review*, 89: 335-339.
- BOLTON, GARY E. AND AXEL OCKENFELS 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *The American Economic Review*, 90: 166-193.
- BOLTON, GARY E. AND RAMI ZWICK 1995. "Anonymity versus Punishment in Ultimatum Bargaining." *Games and Economic Behavior*, 10: 95-121.
- BONJOUR, LAURENCE 1985. *The Structure of Empirical Investigation*. Cambridge, MA: Harvard University Press.
- BOON, MIEKE 2004. "Review: Technological Instruments in Scientific Experimentation". *International Studies in the Philosophy of Science*, 18: 221-30.
- BOSTIC, R., R. J. HERRNSTEIN AND R. D. LUCE 1990. "The Effect on the Preference-Reversal Phenomenon of Using Choice Indifferences". *Journal of Economic Behavior and Organization*, 13: 193-212.
- BOWLES, SAMUEL 1998. "Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions". *Journal of Economic Literature*, 39: 75-111.
- BOWLES, SAMUEL AND HERBERT GINTIS 2004. "The Evolution of Strong Reciprocity: Cooperation in Heterogeneous Populations". *Theoretical Population Biology*, 65:17-28.
- BUCHWALD, JED Z. (ed.) 1995. *Scientific Practice: Theories and Stories of Physics*. Chicago: Chicago University Press.
- BURLANDO, ROBERTO M. AND FRANCESCO GUALA 2005. "Heterogeneous Agents in Public Goods Experiments". *Experimental Economics*, 8: 35-54.
- CALLON, MICHEL 1987. "Society in the Making: The Study of Technology as a Tool for Sociological Analysis". In W. Bijker et al. (eds), pp. 83-103.
- CALLON, MICHEL 1991. "Techno-economic networks and irreversibility". In J. Law, pp. 132-165.
- CALLON, MICHEL AND BRUNO LATOUR 1992. "Don't Throw the Baby Out with the Bath School! A Reply to Collins and Yearley". In A. Pickering, pp. 343-68.
- CAMERER, COLIN F. 1995. "Individual Decision Making". In J.H. Kagel and A. E. Roth (eds.), pp. 587-703.
- CAMERER, COLIN F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, New Jersey: Princeton University Press.
- CAMERER, COLIN F. AND GEORGE LOEWENSTEIN. 2004. "Behavioral Economics: Past, Present, Future". In C. F. Camerer et al. (eds.), pp. 3-51.
- CAMERER, COLIN F., GEORGE LOEWENSTEIN AND MATTHEW RABIN (eds.) 2004. *Advances in Behavioral Economics*. Princeton: Princeton University Press.
- CAMERER, COLIN F. AND RICHARD H. THALER 1995. "Ultimatums, Dictators and Manners." *Journal of Economic Perspectives*, 9: 209-219.
- CAMERER, COLIN F. AND ROBIN M. HOGARTH 1999. "The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework." *Journal of Risk and Uncertainty*, 19: 7-42.
- CAMERON, LISA 1999. "Raising the Stakes in the Ultimatum Game: Experimental Evidence from Indonesia." *Economic Inquiry*, 37: 47-59.
- CAPEN, E. C., R. V. CLAPP AND W. M. CAMPBELL 1971. "Competitive bidding in high-risk situations". *Journal of Petroleum Technology*, 23: 641-53.
- CARNAP, RUDOLF 1967 [1928]. *The Logical Structure of the World*. Berkeley: University of California Press. [Transl. from German]
- CARTWRIGHT, NANCY 1999. *The Dappled World*. Cambridge: Cambridge University Press.
- Chamberlin, Edward H. 1948. "An Experimental Imperfect Market". *Journal of Political Economy*, 56: 95-108.
- CHEW, SOO HONG AND KENNETH MACCRIMMON 1979. "Alpha-nu Choice Theory: a Generalisation of Expected Utility Theory." Working Paper, 669, U. British Columbia.
- CHEW, SOO HONG 1984. "An Axiomatization of the Rank-Dependent Quasilinear Mean Generalizing the Gini Mean and the Quasilinear Mean". Working Papers in Economics, John Hopkins University.
- CHU, YUN-PENG AND RUEY-LING CHU 1990. "The Subsidence of Preference Reversals in Simplified and Market Experimental Settings: A Note". *The American Economic Review*, 80: 902-911.
- COLLINS, HARRY M. 1985. *Changing Order: Replication and Induction in Scientific Practice*. London: Sage Publications.
- COLODNY, ROBERT G. (ed.) 1965. *Beyond the Edge of Certainty*. Prentice-Hall, Englewood Cliffs, NJ.
- COURSEY, DON, R. MARK ISAAC AND VERNON L. SMITH 1984. "Natural Monopoly and the Contested Markets: Some Experimental Results." *Journal of Law and Economics*, 27: 91-113.
- COX, JAMES C. AND DAVID M. GREETHER 1996. "The Preference Reversal Phenomenon: Response Mode, Markets and Incentives". *Economic Theory*, 7: 381-405.
- COX, JAMES C. AND SETH EPSTEIN 1989. "Preference Reversals Without the Independence Axiom". *The American Economic Review*, 79: 408-426.
- CROSS, JOHN G. 1980. "Some Comments on the Papers by Kagel and Battalio and Smith." In J. Kmenta and J. B. Ramsey (eds.), pp. 403-6.
- CUBITT, ROBIN 2005. "Experiments and the Domain of Economic Theory." *Journal of Economic Methodology*, 12: 197-210.
- CUBITT, ROBIN P., CHRIS STARMER AND ROBERT SUGDEN 1998. "On the Validity of the Random Lottery Incentive System". *Experimental Economics*, 1: 115-131.
- CUBITT, ROBIN P., CHRIS STARMER AND ROBERT SUGDEN 2001. "Discovered Preferences and the Experimental Evidence of Violations of Expected Utility Theory". *Journal of Economic Methodology*, 8: 385-414.
- DASGUPTA, PARTHA AND DAVID, PAUL A. 1994. "Toward a new Economics of Science". *Research Policy*, 23: 487-521.
- DAVIS, DOUGLAS AND HOLT, CHARLES. 1993. *Experimental Economics*. Princeton, NJ: Princeton University Press.
- DUHEM, PIERRE. 1954 [1906]. *The Aim and Structure of Physical Theory*. Princeton: Princeton University Press. [original published in French]
- EATWELL, JOHN, MURRAY MILGATE AND PETER NEWMAN (eds.) 1987. *The New Palgrave: A Dictionary of Economic Theory and Doctrine*. Macmillan Press.
- FALK, ARMIN, ERNST FEHR AND URS FISCHBACHER 2003. "On the Nature of Fair Behavior." *Economic Inquiry*, 41: 20-26.
- FEHR, ERNST AND KLAUS M. SCHMIDT 1999. "A Theory of Fairness, Competition, and Cooperation." *The Quarterly Journal of Economics*, 114: 817-868.
- FEHR, ERNST AND SIMON GÄCHTER 2000. "Fairness and Retaliation: The Economics of Reciprocity." *Journal of Economic Perspectives*, 14: 159-181.
- FEYERABEND, PAUL K. [1975] 1993. *Against Method*. 2nd ed. London: Verso.
- FEYERABEND, PAUL K. 1978. *Science in a Free Society*. London: New Left Books.
- FEYERABEND, PAUL K. 1965. "Problems of Empiricism". In R. G. Colodny (ed.), pp. 145-260.
- FIORINA, M. P. AND CHARLES R. PLOTT. 1978. "Committee Decisions Under Majority Rule: An Experimental Study." *American Political Science Review*, 63: 561-604.
- FISHBURN, PETER C. 1982. "Nontransitive Measurable Utility." *Journal of Mathematical Psychology*, 26: 31-67.

- FISHBURN, PETER C. 1983. "Transitive Measurable Utility." *Journal of Economic Theory*, 31: 293-317.
- FLOOD, MERRILL M. 1958. "Some Experimental Games." *Management Science*, 5: 5-26.
- FORSYTHE, ROBERT, JOEL HOROWITZ, N. E. Savin and Martin Sefton 1994. "Fairness in Simple Bargaining Games." *Games and Economic Behavior*, 7: 347-369.
- FOURAKER, LAWRENCE E. AND SIDNEY SIEGEL 1963. *Bargaining Behavior*. New York: McGraw-Hill.
- FRANKLIN, ALLAN 1986. *The Neglect of Experiment*. Cambridge: Cambridge University Press.
- FRANKLIN, ALLAN 1989. "The Epistemology of Experiment." In Gooding et al., pp. 437-60.
- FRANKLIN, ALLAN 1990a. *Experiment, Right or Wrong*. Cambridge: Cambridge University Press.
- FRANKLIN, ALLAN 1991. "Do Mutants Have to Be Slain, or Do They Die of Natural Causes?" In *PSA 1990*, 2: 487-494.
- FRANKLIN, ALLAN 1998. "Experiments in Physics." In E. N. Zalta (ed.) *The Stanford Encyclopaedia of Philosophy*, <http://plato.stanford.edu/entries/physics-experiment>.
- FRANKLIN, ALLAN 1999. *Can that be Right?* Dordrecht: Kluwer Academic Publishers.
- FRIEDMAN, DANIEL AND SHYAM SUNDER. 1994. *Experimental Methods*. Cambridge: Cambridge University Press.
- FRIEDMAN, MICHAEL 1974. "Explanation and Scientific Understanding". *Journal of Philosophy*, 71: 5-19.
- FRIEDMAN, MILTON 1953. *Essays in Positive Economics*. Chicago: The University of Chicago Press.
- FRIEDMAN, MILTON 1953. "The methodology of positive economics" in M. Friedman, pp. 3-43.
- GALISON, PETER 1987. *How Experiments End*. Chicago: The University of Chicago Press.
- GALISON, PETER 1995. "Context and Constraints". In Jed Z. Buchwald (ed.), pp. 13-41.
- GALISON, PETER 1997. *Image and Logic*. Chicago: University of Chicago Press.
- GIERE, RONALD N. 1988. *Explaining Science: A Cognitive Approach*. Chicago: Chicago University Press.
- GIBBARD, ALLAN F. AND HAL R. VARIAN 1978. "Economic Models." *Journal of Philosophy*, 75: 664-677.
- GINTIS, HERBERT 2000. "Strong Reciprocity and Human Sociality." *Journal of Theoretical Biology*, 206: 169-179.
- GINTIS, HERBERT, SAMUEL BOWLES, ROBERT BOYD AND ERNST FEHR 2005. "Moral Sentiments and Material Interests: Origins, Evidence, and Consequences" in H. Gintis et al. (eds), pp. 3-39.
- GINTIS, HERBERT, SAMUEL BOWLES, ROBERT BOYD AND ERNST FEHR (eds.) 2005. *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. Cambridge, Mas.: The MIT Press.
- GODE, DHANANJAY K. AND SHYAM SUNDER 1993. "Allocative Efficiency of Markets with Zero-Intelligence Traders: Markets as a Partial Substitute for Individual Rationality." *Journal of Political Economy*, 101: 119-37.
- GOLDMAN, ALVIN 1978. "Epistemics: The Regulative Theory of Cognition". *The Journal of Philosophy*, 75: 509-523.
- GOLDMAN, ALVIN 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- GOLDMAN, ALVIN 1987. "Foundations of Social Epistemics". *Synthese*, 73: 109-44.
- GOLDMAN, ALVIN AND SHAKED, M. 1991. "An Economic Model of Scientific Activity and Truth Acquisition." *Philosophical Studies* 63: 31-55.
- GOODING, DAVID C. 1985. "In Nature's School: Faraday as an Experimentalist." In D. C. Gooding (ed.), pp. 105-135.
- GOODING, DAVID C. (ed.) 1985. *Faraday Rediscovered: Essays on the Life and Work of Michael Faraday*. Basingstake: Macmillan.
- GOODING, DAVID C. 1986. "How Scientists Reach Agreement about New Observations." In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pp. 236-244.
- GOODING, DAVID C. 1989. "Magnetic Curves and the Magnetic Field: Experimentation and Representation in the History of a Theory." In D. C. Gooding et al. (eds.), pp. 183-223.
- GOODING, DAVID C. 1990. *The Making of Meaning*. Dordrecht: Martinus Nijhoff.
- GOODING, DAVID C. 1992. "Putting Agency Back Into Experiment". In A. Pickering (ed.), pp. 65-112.
- GOODING, DAVID C. 1998. "Picturing Experimental Practice." In M. Heidelberger and F. Steinle (eds.), pp. 298-322.
- GOODING, DAVID C., TREVOR PINCH AND SIMON SCHAFFER (eds.) 1989. *The Uses of Experiment*. Cambridge: Cambridge University Press.
- GRETHER, DAVID M. AND CHARLES R. PLOTT 1979. "Economic Theory of Choice and the Preference Reversal Phenomenon." *The American Economic Review*, 69: 623-638.
- GRETHER, DAVID M. AND CHARLES R. PLOTT 1982. "Economic Theory of Choice and the Preference Reversal Phenomenon: Reply." *The American Economic Review*, 72: 575.
- GRETHER, DAVID M., R. MARK ISAAC AND CHARLES R. PLOTT 1989. *The Allocation of Scarce Resources – Experimental Economics and the Problem of Allocating Airport Slots*. Boulder, San Francisco, London: Westview Press.
- GUALA, FRANCESCO 1998. "Experiments as Mediators in the Non-laboratory Sciences." *Philosophica*, 62: 901-18.
- GUALA, FRANCESCO 1999a. *Economics and the Laboratory*. PhD Dissertation, London School of Economics and Political Science.
- GUALA, FRANCESCO 1999b. "The Problem of External Validity (Or 'Parallelism') in Experimental Economics." *Social Science Information*, 38: 555-73.
- GUALA, FRANCESCO 2000. "Artefacts in Experimental Economics: Preference Reversals and the Becker-DeGroot-Marschak Mechanism." *Economics and Philosophy*, 16: 47-75.
- GUALA, FRANCESCO 2001. "Building Economic Machines: The FCC Auctions." *Studies in History and Philosophy of Science*, 32: 453-77.
- GUALA, FRANCESCO 2002a. "Models, Simulations, and Experiments." In L. Magnani and N.J. Nersessian (eds.), *Model-Based Reasoning: Science, Technology, Values*, New York: Kluwer, p. 59-74.
- GUALA, FRANCESCO 2002b. "On the Scope of Experiments in Economics: Comments on Siakantaris." *Cambridge Journal of Economics*, 26: 261-267.
- GUALA, FRANCESCO 2003. "Experimental Localism and External Validity." *Philosophy of Science*, 70: 1195-205.
- GUALA, FRANCESCO 2005a. "Economics in the Lab: Completeness vs. Testability". *Journal of Economic Methodology*, 12: 185-96.
- GUALA, FRANCESCO 2005b. *The Methodology of Experimental Economics*. New York: Cambridge University Press.
- GÜTH, WERNER, ROLF SCHMITTBERGER AND BERND SCHWARZ 1982. "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization*, 3: 367-88.
- GÜTH, WERNER AND RICHARD TIETZ 1990. "Ultimatum Bargaining Behavior: A Survey and Comparison of Experimental Results." *Journal of Economic Psychology*, 11: 417-49.
- GÜTH, WERNER AND MENACHEM YAARI 1992. "Explaining Reciprocal Behavior in Simple Strategic Games: An Evolutionary Approach." In Ulrich Witt (ed.), pp. 23-34.
- HACKING, IAN 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.
- HACKING, IAN 1989. "Philosophers of Experiment". In *PSA 1988*, pp 147-56.
- HACKING, IAN 1992a. 'The Self-Vindication of the Laboratory Sciences'. In Andrew Pickering (ed.), pp. 29-64.
- HACKING, IAN 1995. "Introduction". In Jed Z. Buchwald (ed.), pp. 1-9.
- HACKING, IAN 1999. *The Social Construction of What?* Cambridge, MA: Harvard University Press.
- HANDS, D. WADE 1993. *Testing, Rationality, and Progress: Essays on the Popperian Tradition in Economic Methodology*. Lanham, MD: Rowman & Littlefield.
- HANDS, D. WADE 1995. "Social Epistemology Meets the Invisible Hand: Kitcher on the Advancement of Science." *Dialogue* 34: 605-21.
- HANDS, D. WADE 2001. *Reflection without Rules: Economic Methodology and Contemporary Science Theory*. Cambridge: Cambridge University Press.
- HANSON, NORWOOD RUSSELL. 1958. *Patterns of Discovery*. Cambridge: Cambridge University Press.
- HARE, R. M. 1973. "Rawls' Theory of Justice". *Philosophical Quarterly*, 23: 144-55; 241-51.

- HARMAN, GILBERT 1965. "The Inference to the Best Explanation." *Philosophical Review*, 74: 88-95.
- HARRÉ, ROM 1970. *The Principles of Scientific Thinking*. London: Macmillan.
- HARRISON, GLENN W. 1994. "Expected Utility Theory and the Experimentalists." *Empirical Economics*, 19: 223-53.
- HAUSMAN, DANIEL M. 1992. *The Inexact and Separate Science of Economics*. Cambridge: Cambridge University Press.
- HAUSMAN, DANIEL M. 2005. "'Testing' Game Theory". *Journal of Economic Methodology*, 12: 211-223.
- HEMPEL, CARL G. AND OPPENHEIM, PAUL. 1948. "Studies in the Logic of Explanation". *Philosophy of Science*, 15: 135-75.
- HEIDELBERGER, MICHAEL AND FRIEDRICH STEINLE (eds.) 1998. *Experimental Essays: Versuche zum Experiment*. Baden-Baden: Nomos Verlagsgesellschaft.
- HENRICH, JOSEPH, ROBERT BOYD, SAMUEL BOWLES, COLIN CAMERER, ERNST FEHR, HERBERT GINTIS AND RICHARD MCELREATH 2001. "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies." *American Economic Review*, 91: 73-78.
- HENRICH, JOSEPH, ROBERT BOYD, SAMUEL BOWLES, COLIN CAMERER, ERNST FEHR AND HERBERT GINTIS (eds.) 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford: Oxford University Press.
- HERTWIG, RALPH AND ANDREAS ORTMANN 2001. "Experimental Practices in Economics: a Methodological Challenge for Psychologists?" *Behavioral and Brain Sciences*, 24: 383-451.
- HESSE, MARY 1966. *Models and Analogies in Science*. Notre Dame, Ind.: University of Notre Dame Press.
- HEY, JOHN D. 1991. *Experiments in Economics*. Oxford: Basil Blackwell.
- HOFFMAN, ELIZABETH, KEVIN MCCABE AND VERNON L. SMITH 1995. "Ultimatum and Dictator Games." *Journal of Economic Perspectives*, 9: 236-9.
- HOFFMAN, ELIZABETH, KEVIN MCCABE AND VERNON L. SMITH 1996a. "Social Distance and Other-Regarding Behavior in Dictator Games." *The American Economic Review*, 86: 653-60.
- HOFFMAN, ELIZABETH, KEVIN MCCABE AND VERNON L. SMITH 1996b. "On Expectations and Monetary Stakes in Ultimatum Games." *International Journal of Game Theory*, 25: 289-301.
- HOFFMAN, ELIZABETH, KEVIN MCCABE AND VERNON L. SMITH 1999. "Social Distance and Other-Regarding Behavior in Dictator Games: Reply." *The American Economic Review*, 89: 340-1.
- HOFFMAN, ELIZABETH, KEVIN MCCABE, SHACHAT, KEITH AND VERNON L. SMITH 1994. "Preferences, Property Rights and Anonymity in Bargaining Games." *Games and Economic Behavior*, 7: 346-80.
- HOGGATT, AUSTIN C. 1959 "An Experimental Business Game." *Behavioral Science*, 4: 192-203.
- HOLT, CHARLES A. 1986. "Preference Reversals and the Independence Axiom." *The American Economic Review*, 76: 508-515.
- HOLT, CHARLES A. 1995. "Industrial Organization: A Survey of Laboratory Research." In J. H. Kagel and A. E. Roth (eds.), pp 349-443.
- HOLT, CHARLES A., LORIN LANGAN AND ANNE VILLAMIL 1986. "Market Power in Oral Double Auction." *Economic Inquiry*, 24: 107-23.
- HON, GIORA 1989. "Towards a Typology of Experimental Errors: An Epistemological View". *Studies in History and Philosophy of Science* 20, pp. 469-504.
- HON, GIORA 2003. "The Idols of Experiment: Transcending the 'Etc. List'". In Hans Radder (ed.), pp. 174-98.
- HURWICZ, LEONID 1960. "Optimality and Informational Efficiency in Resource Allocation Processes." In K. Arrow et al. (eds.), pp. 27-46.
- ISAAC, R. MARK AND VERNON L. SMITH 1985. "In Search of Predatory Pricing." *Journal of Political Economy*, 93: 320-45.
- JANICH, PETER 1978. "Physics – Natural Science or Technology?" In W. Krohn et al. (eds.), pp 3-27.
- KAGEL, JOHN H. AND ALVIN E. ROTH (eds.) 1995. *The Handbook of Experimental Economics*, Princeton: Princeton University Press.
- KAGEL, JOHN H. 1995. "Auctions: A Survey of Experimental Research." In J. H. Kagel and A. E. Roth (eds.), pp. 501-85.
- KAGEL, JOHN H. AND DAN LEVIN 1986. "The Winner's Curse and Public Information in Common Value Auctions." *American Economic Review*, 76: 894-920.
- KAHNEMAN, DANIEL AND AMOS TVERSKY 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, 47: 263-91.
- KAHNEMAN, DANIEL, JACK L. KNETSCH AND RICHARD THALER 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *The American Economic Review*, 76: 728-41.
- KALISCH, GERHARD K., J. W. MILNOR, JOHN F. NASH AND E. D. NERING 1954. "Some experimental N-persons Games." In R. M. Thrall et al. (eds.), pp. 301-327.
- KARNI, EDI, AND ZVI SAFRA 1987. "Preference Reversal" and the Observability of Preferences by Experimental Methods." *Econometrica*, 55: 675-85.
- KELLER, L. R., UZI SEGAL AND T. WANG 1993. "The Becker-DeGroot-Marschak Mechanism and Generalized Utility Theories: Theoretical Predictions and Empirical Observations." *Theory and Decision*, 34: 83-97.
- KITCHER, PHILIP 1990. "The Division of Cognitive Labor". *The Journal of Philosophy*, 87: 5-22.
- KITCHER, PHILIP 1993. *The Advancement of Science*. Oxford: Oxford University Press.
- KITCHER, PHILIP 1994. "Contrasting Conceptions of Social Epistemology." In F. Schmitt (ed.), pp: 111-34.
- KITCHER, PHILIP 2001. *Science, Truth, and Democracy*. New York: Oxford University Press.
- KITCHER, PHILIP 2002. "Reply to Helen Longino." *Philosophy of Science*, 69: 569-72.
- KMENTA, JAN AND JAMES B. RAMSEY (eds.) 1980. *Evaluation of Econometric Models*. Academic Press.
- KNEZ, MARC AND VERNON L. SMITH 1987. "Hypothetical valuations and preference reversals in the context of asset trading." In A. E. Roth (ed.), pp. 131-154.
- KNORR-CETINA, KARIN 1981. *The Manufacture of Knowledge: An Essay on the Constructivist and Contextual Nature of Science*. New York: Pergamon.
- KROHN, WOLFGANG, EDWIN T. LAYTON AND PETER WEINGART (eds.) 1978. *The Dynamics of Science and Technology*, Dordrecht, Boston: D. Reidel Publishing Company.
- KUHN, THOMAS S. 1970 [1962]. *The Structure of Scientific Revolutions*, 2nd ed. Chicago: University of Chicago Press.
- KUHN, THOMAS S. 1977. *The Essential Tension*. Chicago: University of Chicago Press.
- Lakatos, Imre. 1970. "Falsification and the Methodology of Scientific Research Programmes." In I. Lakatos and A. Musgrave (eds.), pp. 91-196.
- LAKATOS, IMRE AND ALAN MUSGRAVE (eds.) 1970. *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.
- LATOUR, BRUNO. 1987. *Science in Action*. Cambridge, Mass.: Harvard University Press.
- Latour, Bruno and Steven Woolgar 1986 [1979]. *Laboratory Life: The Construction of Scientific Facts*. Princeton: Princeton University Press, 2nd ed.
- LAUDAN, LARRY 1977. *Progress and Its Problems: Towards a Theory of Scientific Growth*. Berkeley: University of California Press.
- LAW, JOHN (ed.) 1991. *A Sociology of Monsters: Essays on Power, Technology and Domination*. London: Routledge.
- LEAMER, EDWARD 1983. "Let's Take the Con Out of Econometrics". *American Economic Review*, 73: 31-64.
- LESTER, R. A. 1946. "Shortcomings of Marginal Analysis for Wage-Employment Problems". *American Economic Review*, 36: 62-82.
- LESTER, R. A. 1947. "Marginalism, Minimum Wages, and Labour Markets". *American Economic Review*, 37: 135-48.
- LEVINE, DAVID K. 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics*, 1: 593-622.
- LICHTENSTEIN, SARAH AND PAUL SLOVIC 1971. "Reversals of Preference Between Bids and Choices in Gambling Decisions." *Journal of Experimental Psychology*, 89: 46-55.

- LICHTENSTEIN, SARAH AND PAUL SLOVIC 1973. "Response-Induced Reversals of Preference in Gambling: An Extended Replication in Las Vegas." *Journal of Experimental Psychology*, 101: 16-20.
- LINDMAN, H. R. 1971. "Inconsistent Preferences among Gambles." *Journal of Experimental Psychology*, 89: 390-7.
- LIPSEY, RICHARD 1979. *An Introduction to Positive Economics*. 5th ed. London: Weidenfeld and Nicolson.
- LIPSEY, RICHARD AND ALEC CRISTAL 1995. *An Introduction to Positive Economics*. 8th ed. Oxford: Oxford University Press.
- LOEWENSTEIN, GEORGE 1999. "Experimental Economics from the Vantage-Point of Behavioural Economics." *The Economic Journal*, 109: F25-34.
- LONGINO, HELEN E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton, N.J.: Princeton University Press.
- LONGINO, HELEN E. 1991. "Multiplying Subjects and the Diffusion of Power." *The Journal of Philosophy* 8: 666-74.
- LONGINO, HELEN E. 1994. "The Fate of Knowledge in Social Theories of Science." In Frederick F. Schmitt (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*. p. 135-57.
- LONGINO, HELEN E. 2002a. *The Fate of Knowledge*. Princeton: Princeton University Press.
- LONGINO, HELEN E. 2002b. "Science and the Common Good: Thoughts on Philip Kitcher's *Science, Truth, and Democracy*." *Philosophy of Science*, 69: 560-68.
- LOOMES, GRAHAM 1991. "Experimental Methods in Economics." In D. Greenaway, M. Bleaney and I. Stewart (eds.), *Companion to Contemporary Economic Thought*, London: Routledge, p. 593-613.
- LOOMES, GRAHAM 1998. "Probabilities vs Money: A Test of Some Fundamental Assumptions about Rational Decision Making." *The Economic Journal*, 108: 477-89.
- LOOMES, GRAHAM 1999a. "Experimental Economics: Introduction." *The Economic Journal*, 109: F1-4.
- LOOMES, GRAHAM 1999b. "Some Lessons from Past Experiments and Some Challenges for the Future." *The Economic Journal*, 109: F35-45.
- LOOMES, GRAHAM AND ROBERT SUGDEN 1982. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty." *The Economic Journal*, 92: 805-24.
- Loomes, Graham and Robert Sugden 1983. "A Rationale for Preference Reversal." *The American Economic Review*. 73: 428-32.
- LOOMES, GRAHAM, CHRIS STARMER AND ROBERT SUGDEN 1989. "Preference Reversal: Information-Processing Effect or Rational Non-Transitive Choice?" *The Economic Journal*, 99: 140-51.
- LOOMES, GRAHAM, CHRIS STARMER AND ROBERT SUGDEN 1991. "Observing Violations of Transitivity by Experimental Methods." *Econometrica*, 59: 425-39.
- LOPES, HELENA, JOSÉ M. C. CALDAS, ANA C. COSTA, LUÍS F. CARVALHO, J. G. GUSMÃO, JOÃO LEÃO, JOÃO RODRIGUES AND ANA C. SANTOS 2004. "Public Good Provision – Why People do (not) Contribute? An Experimental Contribution". Working paper Dinâmia 36.
- MÄKI, USKALI 1992. "On the Method of Isolation in Economics". *Poznan Studies in the Philosophy of the Sciences and the Humanities*, 26: 319-354.
- MÄKI, USKALI 1993. "Social Theories of Science and the Fate of Institutionalism in Economics". In U. Mäki et al. (eds.), pp. 76-109.
- MÄKI, USKALI 1995. "Diagnosing McCloskey". *Journal of Economic Literature*, 33: 1300-18.
- MÄKI, USKALI 1996. "Scientific Realism and Some Peculiarities of Economics". *Boston Studies in the Philosophy of Science*, 169: 425-445.
- MÄKI, USKALI 1998. "Aspects of Realism about Economics". *THEORIA*, 13: 301-19.
- MÄKI, USKALI 1999. "Sciences as a Free Market: a Reflexivity Test in an Economics of Economics". *Perspectives on Science*, 7: 486-509.
- MÄKI, USKALI 2001. "Models". In *International Encyclopedia of the Social and Behavioral Sciences*, Volume 15, pp. 9931-7. Elsevier.
- MÄKI, USKALI 2002. "Some Non-reasons for Non-realism about Economics". In U. Mäki (ed.), pp. 90-104.
- MÄKI, USKALI (ed.) 2002. *Fact and Fiction in Economics. Realism, Models and Social Construction*. Cambridge: Cambridge University Press.
- MÄKI, USKALI 2005. "Models are Experiments, Experiments are Models". *Journal of Economic Methodology*, 12: 303-15.
- MÄKI, USKALI AND BO GUSTAFSSON AND CHRISTIAN KNUDSEN (eds.) 1993. *Rationality, Institutions and Economic Methodology*. London: Routledge.
- MACHINA, MARK J. 1982. "Expected Utility' Analysis Without the Independence Axiom." *Econometrica*, 50: 277-323.
- MACHLUP, FRITZ 1946. "Marginal Analysis and Empirical Research". *American Economic Review*, 36: 519-54.
- MACHLUP, FRITZ 1947. "Rejoinder to an Antimarginalist". *American Economic Review*, 37: 148-54.
- MACKENZIE, DONALD 1981. "Interests, Positivism and History." *Social Studies of Science*, 11: 498-504.
- MAGNANI, LORENZO AND NANCY J. NERSESSIAN (eds.) 2002. *Model-Based Reasoning: Science, Technology, Values*. Dordrecht: Kluwer.
- MASTERMAN, MARGARET 1970. "The Nature of a Paradigm." In I. Lakatos and A. Musgrave, pp. 59-89.
- MAYER, THOMAS 1980. "Economics as Hard Science: Realistic Goal or Wishful Thinking". *Economic Inquiry*, 18:165-177.
- MAYO, DEBORAH G. 1996. *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.
- MCCABE, KEVIN, STEPHEN J. RASSENTI AND VERNON L. SMITH 1989. "Designing 'Smart' Computer-Assisted Markets: An Experimental Auction for Gas Networks." *European Journal of Political Economy*, 5: 259-83.
- MCCLOSKEY, D. 1985. *The Rhetoric of Economics*. Madison: University of Wisconsin Press.
- MCDANIEL, TANGA AND CHRIS STARMER 1998. "Experimental Economics and Deception: A Comment." *Journal of Economic Psychology*, 19: 403-9.
- MILGROM, PAUL R. AND ROBERT J. WEBER 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50: 1089-122.
- MILL, JOHN STUART 1884 [1843]. *A System of Logic*. 8th ed. New York: Harper & Brothers.
- MILL, JOHN STUART. 1874 [1844]. "On the Definition and Method of Political Economy; and the Method of Investigation Proper to It". In *Essays on Some Unsettled Questions of Political Economy*, 2nd ed., London: Longmans, Green, Reader & Dyer.
- MILL, JOHN STUART. 1982 [1859]. *On Liberty*. Harmondsworth: Penguin.
- MILLER, ROSS M., CHARLES R. PLOTT AND VERNON L. SMITH 1977. "Intertemporal Competitive Equilibrium: An Empirical Study of Speculation." *Quarterly Journal of Economics*, 91: 599-624.
- MIROWSKI, PHILIP 1996. "The Economic Consequences of Philip Kitcher." *Social Epistemology* 10: 153-69.
- MIROWSKI, PHILIP 2002. *Machine Dreams: Economics Becomes a Cyborg Science*. Cambridge: Cambridge University Press.
- MORGAN, MARY 2002. "Model Experiments and Models in Experiments." In L. Magnani and N.J. Nersessian (eds.), pp. 41-58.
- MORGAN, MARY 2003. "Experiments without Material Intervention: Model Experiments, Virtual experiments and Virtually Experiments." In H. Radder (ed.), pp. 216-35.
- MORGAN, MARY 2005. "Experiments versus Models: New Phenomena, Inference and Surprise". *Journal of Economic Methodology*, 12: 317-29.
- MORGAN, MARY S. AND MORRISON, MARGARET C. (eds.) 1999. *Models as Mediators*. Cambridge: Cambridge University Press.
- MORRISON, MARGARET C. AND MARY S. MORGAN 1999. "Models as Mediating Instruments." In M. S. Morgan and M. C. Morrison (eds.), pp 10-37.
- MULKAY, M. AND GILBERT, G. N. 1986. "Replication and Mere Replication." *Philosophy of the Social Sciences*, 16: 21-37.

- NERSESSIAN, NANCY J. (ed.) 1987. *The Process of Science*. Dordrecht: Martinus Nijhoff Publishers.
- NEWTON-SMITH, WILLIAM H. 1981. *The Rationality of Science*. London: Routledge & Kegan Paul.
- NICKLES, THOMAS 1989. "Justification and Experiment" in D. Gooding et al., pp. 299-333.
- Ochs, Jack and Alvin E. Roth 1989. "An Experimental Study of Sequential Bargaining." *The American Economic Review*, 79: 355-84.
- PALFREY, THOMAS AND ROBERT PORTER 1991. "Guidelines for Submission of Manuscripts on Experimental Economics." *Econometrica*, 59: 1197-98.
- PICKERING, ANDREW 1981. "The Hunting of the Quark". *Isis*, 72: 216-36.
- PICKERING, ANDREW 1984. *Constructing Quarks*. Chicago: University of Chicago Press.
- PICKERING, ANDREW 1989. "Living in the Material World: On Realism and Experimental Practice." In D. C. Gooding et al. (eds.), pp. 275-97.
- PICKERING, ANDREW 1990a. "Reason Enough? More on Parity-Violation Experiments and Electroweak Gauge Theory." In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pp. 459-69.
- PICKERING, ANDREW 1990b. "Knowledge, Practice and Mere Construction." *Social Studies of Science*, 20: 682-729.
- PICKERING, ANDREW 1992. "From Science as Knowledge to Science as Practice." In A. Pickering (ed.), pp. 1-26.
- PICKERING, ANDREW (ed.) 1992. *Science as Practice and Culture*, Chicago: University of Chicago Press.
- PICKERING, ANDREW 1994a. "The Mangle of Practice: Agency and Emergence in the Sociology of Science." *American Journal of Sociology*, 99: 559-89.
- PICKERING, ANDREW 1994b. "After Representation: Science Studies in the Performative Idiom." *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, pp. 413-9.
- PICKERING, ANDREW 1995a. "Context and Constraints." In J. Z. Buchwald (ed.), pp. 13-41.
- PICKERING, ANDREW 1995b. "Beyond Constraint: The Temporality of Practice and the Historicity of Knowledge." In J. Z. Buchwald (ed.), pp. 42-55.
- PICKERING, ANDREW 1995c. *The Mangle of Practice: Time, Agency, and Science*. Chicago: University of Chicago Press.
- PINCH, TREVOR 1986. *Confronting Nature*. Dordrecht: Reidel.
- PLOTT, CHARLES R. 1979. "The Application of Laboratory Experimental Methods to Public Choice." In Clifford S. Russell (ed.), *Collective Decision Making: Applications from Public Choice Theory*, Baltimore MD: John Hopkins University Press. p. 137-60.
- PLOTT, CHARLES R. 1982. "Industrial Organization Theory and Experimental Economics". *Journal of Economic Literature*, 20: 1485-527.
- PLOTT, CHARLES R. 1987. "Dimensions of Parallelism: Some Policy Applications of Experimental Methods". In A. E. Roth (ed.), pp. 193-219.
- PLOTT, CHARLES R. 1991. "Will Economics Become an Experimental Science?" *Southern Economic Journal*, 57: 901-19.
- PLOTT, CHARLES R. 1995. "Rational Individual Behaviour in Markets and Social Choice Processes: the Discovered Preference Hypothesis." In K. J. Arrow et al., pp. 225-50.
- PLOTT, CHARLES R. 1997. "Laboratory Experimental Testbeds: Application to the PCS Auction." *Journal of Economics and Management Strategy*, 6: 605-38.
- PLOTT, CHARLES R. AND MICHAEL E. LEVINE. 1978. "A Model Agenda Influence on Committee Decisions." *The American Economic Review*, 88: 146-60.
- PLOTT, CHARLES, AND VERNON L. SMITH. 1978. "An Experimental Examination of Two Exchange Institutions." *Review of Economic Studies*. 45: 133-53
- POLANYI, MICHAEL. 1958. *Personal Knowledge*. Chicago: Chicago University Press.
- POMMERHNE, WERNER W., FRIEDRICH SCHNEIDER AND PETER ZWEIFEL 1982. "Economic Theory of Choice and the Preference Reversal Phenomenon: A Reexamination." *The American Economic Review*, 72: 569-74.
- POPPER, R. KARL 1959. *The Logic of Scientific Discovery*. New York: Basic Books (translation of Popper 1934).
- POPPER, R. KARL 1965 [1963]. *Conjectures and Refutations*. New York: Harper and Row. 2nd ed.
- POPPER, R. KARL 1979 [1972] *Objective Knowledge: An Evolutionary Approach*. Oxford: Clarendon Press. 2nd ed.?
- POPPER, R. KARL 1994. *The Myth of the Framework: In Defense of Science and Rationality*. London: Routledge.
- PRASNIKAR, VESNA AND ALVIN E. ROTH 1992. "Considerations of Fairness and Strategy: Experimental Data from Sequential Games." *The Quarterly Journal of Economics*, 107: 865-88.
- Quiggin, John 1982. "A Theory of Anticipated Utility." *Journal of Economic Behavior and Organization*, 3: 324-43.
- QUINE, WILLARD ORMAN 1951. "Two Dogmas of Empiricism". *Philosophical Review*, 60: 20-43.
- RABIN, MATTHEW 1993. "Incorporating Fairness into Game Theory and Economics." *The American Economic Review*, 83: 1281-302.
- RABIN, MATTHEW 1998. "Psychology and Economics." *Journal of Economic Literature*, 36: 11-46.
- RADDER, HANS 1988. *The Materialization of Science*. Assen: Van Gorcum.
- RADDER, HANS 1995. "Experimenting in the Natural Sciences: A Philosophical Approach." In Jed Z. Buchwald (ed.), *Scientific Practice: Theories and Stories of Physics*, Chicago: Chicago University Press, p. 56-86.
- RADDER, HANS 1996. *In and About the World - Philosophical Studies of Science and Technology*. Albany, New York: State University of New York Press.
- RADDER, HANS 1998. "Issues for a Well-Developed Philosophy of scientific Experimentation." In M. Heidelberger and F. Steinle (eds.), *Experimental Essays: Versuche zum Experiment*, Baden-Baden: Nomos Verlagsgesellschaft, p. 392-404.
- RADDER, HANS (ed.) 2003. *The Philosophy of Scientific Experimentation*. Pittsburgh: University of Pittsburgh Press.
- RADDER, HANS 2003a. "Toward a More Developed Philosophy of Scientific Experimentation". In Radder (ed.), *The Philosophy of Scientific Experimentation*. Pittsburgh: University of Pittsburgh Press, p. 1-18.
- RADDER, HANS 2003b. "Technology and Theory in Experiment." In Radder (ed.), *The Philosophy of Scientific Experimentation*. Pittsburgh: University of Pittsburgh Press, p. 152-73.
- RASSENTI, STEPHEN. J., VERNON L. SMITH AND R. L. BULFIN 1982. "A Combinatorial Auction Mechanism for Airport Time Slot Allocation." *Bell Journal of Economics*, 13: 402-17.
- RAWLS, JOHN 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- RAWLS, JOHN 1999. *A Theory of Justice*. 2nd ed. Cambridge, MA: Harvard University Press.
- REILLY, ROBERT J. 1982. "Preference Reversal: Further Evidence and Some Suggested Modifications in Experimental Design." *The American Economic Review*, 72: 576-84.
- REITER, STANLEY 1977. "Information and Performance in the (New) Welfare Economics." *American Economic Review Proceedings*, 67: 226-34.
- ROBBINS, LIONEL 1932. *An Essay on the Nature and Significance of Economic Science*. London: Macmillan.
- ROTH, ALVIN E. (ed.) 1987. *Laboratory Experimentation in Economics: Six Points of View*. Cambridge: Cambridge University Press.
- ROTH, ALVIN E. 1988. "Laboratory Experimentation in Economics: A Methodological Overview." *Economic Journal*, 98: 974-1031.
- ROTH, ALVIN E. 1995a. Introduction to Experimental Economics. In J. H. Kagel and A. E. Roth (eds.), pp. 3-109.
- ROTH, ALVIN E. 1995b. "Bargaining Experiments." In J. H. Kagel and A. E. Roth (eds.), pp. 253-348.
- ROTH, ALVIN E. AND IDO EREV 1995. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior*, 8: 164-212.
- ROTH, ALVIN E., VESNA PRASNIKAR, MASAHIRO OKUNO-FUJIWARA AND SHMUEL ZAMIR 1991. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *The American Economic Review*, 81: 1068-95.
- SAMUELSON, LARRY 2001. "Adaptations, Analogies and Anomalies." *Journal of Economic*

- Theory*, 97: 320-66.
- SAMUELSON, PAUL A. AND WILLIAM D. NORDHAUS 1985. 12th ed. *Principles of Economics*. McGraw-Hill, Irwin.
- SAUERMAN, HEINZ, AND RICHARD SELTEN 1960. "An Experiment in Oligopoly." In L. von Bertalanffy and A. Rappoport (eds.), *General Systems Yearbook of the Society for General Systems Research*, Ann Arbor: Society for General Systems Research.
- SHELLING, THOMAS C. 1957. "Bargaining, Communication, and Limited War." *Journal of Conflict Resolution*, 1: 19-36.
- SCHKADE, D. A. AND E. J. JOHNSON 1989. "Cognitive Processes in Preference Reversals." *Organizational Behavior and Human Decision Processes*, 44: 203-31.
- SCHMITT, FREDERICK 1994. "Socializing Epistemology: An Introduction through Two Sample Issues". In F. Schmitt (ed.), pp. 1-27.
- SCHMITT, FREDERICK F. (ed.) 1994. *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham, MD: Roman and Littlefield.
- SCHOTTER, ANDREW, AVI WEISS AND INIGO ZAPATER 1996 "Fairness and Survival in Ultimatum and Dictatorship Games." *Journal of Economic Behavior and Organization*, 31: 37-56.
- Schotter, Andrew, Keith Weigelt and Charles Wilson 1994. "A Laboratory Investigation of Multiperson Rationality and Presentation Effects." *Games and Economic Behavior*, 6: 445-68.
- SCHRAM, ARTHUR 2005. "Artificiality: The Tension Between Internal and External Validity in Economics Experiments." *Journal of Economic Methodology*, 12: 225-37.
- SEGAL, UZI 1988. "Does the Preference Reversal Phenomenon Necessarily Contradict the Independence Axiom?" *The American Economic Review*, 78: 233-36.
- SHAPIN, STEVEN 1982. "History of Science and Its Sociological Reconstructions". *History of Science*, 20: 157-211.
- SIKANTARIS, NIKOS 2000. "Experimental Economics under the Microscope." *Cambridge Journal of Economics*, 24: 267-81.
- SIEGEL, SIDNEY 1961. "Decision Making and Learning Under Varying Conditions of Reinforcement." *Annals of the New York Academy of Science*, 89.
- SIEGEL, SIDNEY AND LAWRENCE FOURAKER 1960. *Bargaining and Group Decision Making*. New York.
- SLONIM, ROBERT AND ALVIN E. ROTH 1998. "Financial Incentives and Learning in Ultimatum and Market Games: An Experiment in the Slovak Republic." *Econometrica*, 66: 569-96.
- SLOVIC, PAUL 1975. "Choice Between Equally Valued Alternatives." *Journal of Experimental Psychology: Human Perception and Performance*, 1: 280-7.
- SLOVIC, PAUL 1995. "The Construction of Preferences." *American Psychologist*, 50: 364-71.
- SLOVIC, PAUL AND SARAH LICHTENSTEIN 1968. "The Relative Importance of Probabilities and Payoffs in Risk Taking." *Journal of Experimental Psychology*, 78: 1-18.
- SLOVIC, PAUL AND SARAH LICHTENSTEIN 1983. "Preference Reversals: A Broader Perspective." *The American Economic Review*, 73: 596-605.
- SMITH, VERNON L. 1962. "An Experimental Study of Competitive Market Behaviour." *The Journal of Political Economy*, 70: 322-3.
- SMITH, VERNON L. 1964. "Effect of Market Organization on Competitive Equilibrium." *Quarterly Journal of Economics*, 78: 181-201.
- SMITH, VERNON L. 1965. "Experimental Auction Markets and the Walrasian Hypothesis." *Journal of Political Economy*, 73: 387-93.
- SMITH, VERNON L. 1967. "Experimental Studies of Discrimination versus Competition in Sealed-Bid Auction Markets." *Journal of Business*, 40: 56-84.
- SMITH, VERNON L. 1976a. "Experimental Economics: Induced Value Theory." *American Economic Review*, 66: 274-9.
- SMITH, VERNON L. 1976b. "Bidding and Auctioning Institution: Experimental Results." In Y. Amihud (ed.), *Bidding and Auctioning for Procurement and Allocation*, New York: New York University Press. p?
- SMITH, VERNON L. 1979a. "An Experimental Comparison of Three Public Good Decision Mechanism." *Scandinavian Journal of Economics*, 81: 198-215.
- SMITH, VERNON L. 1979b. "Incentive Compatible Experimental Processes for the Provision of Public Goods." In V. L. Smith (ed.), *Research in Experimental Economics*, vol.1. Greenwich, Conn.: JAI Press. p?
- SMITH, VERNON L. 1980a. "Experiments with a Decentralized Mechanism for Public Goods Decision." *American Economic Review*, 70: 584-99.
- SMITH, VERNON L. 1980b. "Relevance of Laboratory Experiments to Testing Resource Allocation Theory." In Jan Kmenta and James B. Ramsey (eds.), *Evaluation of Econometric Models*, New York: Academic Press, p.
- SMITH, VERNON L. 1981a. An Empirical Study of Decentralized Institutions of Monopoly Restraint. In J. Quirk and G. Horwich (eds.), *Essays in Contemporary Fields of Economics in Honor of E. T. Weiler (1914-1979)*, West Lafayette: Purdue University Press, p. 83-106.
- SMITH, VERNON L. 1981b. Theory, Experiment and Antitrust Policy. In S. Salop (ed.), *Strategy, Predation, and Antitrust Analysis*, Washington, D.C.: Federal Trade Commission, Bureau of Economics, p. 579-603.
- SMITH, VERNON L. 1982. "Microeconomic Systems as an Experimental Science." *American Economic Review*, 72: 923-55.
- SMITH, VERNON L. (ed.) 1985. *Research in Experimental Economics*, Vol. 3, Greenwich, CT: JAI Press.
- SMITH, VERNON L. 1987. "Experimental Methods in Economics". In J. Eatwell et al., pp. 241-8.
- SMITH, VERNON L. 1989. "Theory, Experiment and Economics". *Journal of Economic Perspectives*, 3: 151-69.
- SMITH, VERNON L. 1990. *Experimental Economics. Schools of Thought in Economics*, Vol. 7 Aldershot, England: Edward Elgar; and Brookfield, V.: Gower Publishing.
- SMITH, VERNON L. 1991. "Experimental Economics at Purdue." In Vernon L. Smith (ed.), *Papers in Experimental Economics*, Cambridge: Cambridge University Press. p.
- SMITH, VERNON L. 1992. "Game Theory and Experimental Economics: Beginnings and Early Influences." In E. R. Weintraub (ed.), *Toward a History of Game Theory*, Annual Supplement to vol. 24, History of Political Economy, Durham, N.C.: Duke University Press, p. 241-82.
- SMITH, VERNON L. 1994. "Economics in the Laboratory." *Journal of Economic Perspective*, 8: 113-31.
- SMITH, VERNON L. 2002. "Method in Experiment: Rhetoric and Reality". *Experimental Economics*, 5: 91-110.
- SMITH, VERNON L. AND A. W. WILLIAMS. 1983. "An Experimental Comparison of Alternative Rules for Competitive Market Exchange." In R. Englebrecht-Williams et al. (eds.), *Auctions, Bidding and Controlling: Uses and Theory*, New York University Press.
- SMITH, VERNON L., A. W. WILLIAMS, W.K. BRATTON AND M.G. VANNONI. 1982. "Double Auctions vs Sealed Bid-offer Auction." *American Economic Review*, 72: 58-77.
- SMITH, VERNON L., GERRY L. SUCHANEK AND ARLINGTON W. WILLIAMS. 1988. "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets." *Econometrica*, 56: 1119-52.
- SMITH, VERNON L., K MCCABE AND S. RASSENTI. 1991. "Lakatos and Experimental Economics." In N. de Marchi and M. Blaug (eds.), *Appraising Economic Theories*, Edward Elgar.
- SOLOMON, MIRIAM 1992. "Scientific Rationality and Human Reasoning". *Philosophy of Science*, 59: 439-54.
- SOLOMON, MIRIAM 1994a. "Social Empiricism". *Noûs*, 28: 325-43.
- SOLOMON, MIRIAM 1994b. "A More Social Epistemology." In Frederick F. Schmitt (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*, p. 217-33, Lanham, Md: Rowman & Littlefield.
- SOLOMON, MIRIAM 1995. "Legend Naturalism and Scientific Progress." *Studies in History and Philosophy of Science*, 26: 205-18.
- SOLOMON, MIRIAM 2001. *Social Empiricism*. Cambridge, Mass.: MIT Press.
- Starmer, Chris 1999a. "Experimental Economics: Hard Science or Wasteful Tinkering?" *The Economic Journal*, 109: F5-15.
- STARMER, CHRIS 1999b. "Experiments in Economics: Should we Trust the Dismal Scientists in White Coats?" *Journal of Economic Methodology*, 6: 1-30
- STARMER, CHRIS 2000. "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk." *Journal of Economic Literature*, 38: 332-82.
- STARMER, CHRIS AND ROBERT SUGDEN 1991. "Does the Random-lottery Incentive System

- Elicit True Preference? An Experimental Investigation." *The American Economic Review*, 81: 971-8.
- STARMER, CHRIS AND ROBERT SUGDEN 1998. "Testing Alternative Explanations of Cyclical Choices." *Economica*, 65: 347-61.
- SUGDEN, ROBERT 2002. "Credible Worlds. The Status of Theoretical Models in Economics." In Uskali Mäki (ed.), *Fact and Fiction in Economics: Models, Realism, and Social Construction*, Cambridge: Cambridge University Press, p. 107-36.
- SUGDEN, ROBERT 2005. "Experiments as Exhibits and Experiments as Tests." *Journal of Economic Methodology*, 12: 291-302.
- SUNDER, SHYAM. 1995. "Experimental Asset Markets: A Survey." In J. H. Kagel and A. E. Roth (eds.), *The Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press. p.?
- SUPPES, PATRIK 1962. "Models of Data." *Logic, Methodology and Philosophy of Science*. Stanford: Stanford University Press. p. 252-261. ???
- TAMMI, TIMO 1999a. "Incentives and Preference Reversals: Escape Moves and Community Decisions in Experimental Economics." *Journal of Economic Methodology*, 6: 351-80.
- TAMMI, TIMO 1999b. "Simon's and Siegel's Responses to the 'Mixed Strategy Anomaly': A Missed Case in the Sensitivity of Economics to Empirical Evidence". *Cambridge Journal of Economics*, 27: 85-96.
- TAMMI, TIMO 2003. "On Experimental Discourse in Economics". *Philosophy of the Social Sciences*, 29: 62-88.
- THAGARD, PAUL 1989. "Explanatory Coherence". *Behavioral and Brain Sciences*, 12: 435-467.
- THAGARD, PAUL 1992. *Conceptual Revolutions*. Princeton: Princeton University Press.
- THAGARD, PAUL 1996. *Mind: Introduction to Cognitive Science*. Cambridge MA: MIT Press.
- THALER, RICHARD H. 1988a. "The Ultimatum Game." *Journal of Economic Perspectives*, 2: 195-206.
- THALER, RICHARD H. 1988b. The Winner's Curse. *Journal of Economic Perspectives*, 2: 191-202.
- THRALL, R. M., C. H. COOMBS AND R. L. DAVIS (eds.). 1954. *Decision Processes*. New York: Wiley.
- TVERSKY, AMOS AND RICHARD H. THALER. 1990. "Anomalies: Preference Reversals." *Journal of Economic Perspectives*, 4: 201-11.
- TVERSKY, AMOS, PAUL SLOVIC AND DANIEL KAHNEMAN. 1990. "The Causes of Preference Reversal." *The American Economic Review*, 80: 204-17.
- TVERSKY, AMOS, PAUL SLOVIC AND SHMUEL SATTAH. 1988. "Contingent Weighting in Judgment and Choice." *Psychological Review*, 95: 371-84.
- VAN FRAASSEN, BAS. 1980. *The Scientific Image*. Oxford: Clarendon Press.
- VICKREY, WILLIAM. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16: 8-37.
- VON NEUMANN, JOHN AND OSKAR MORGENSTERN. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- WHEWELL, WILLIAM 1967 [1840]. *The Philosophy of Inductive Sciences*. Johnson Reprint.
- WILDE, LOUIS L. 1981. "On the use of laboratory Experiments in Economics". In J.C. Pitt (ed.), *The Philosophy of Economics*, Dordrecht, Reidel. p. 137-48.
- WILSON, ROBERT 1977. "A Bidding Model of Perfect Competition." *Review of Economic Studies*, 4: 511-8.
- WITT, ULRICH (ed.) 1992. *Explaining Process and Change: Approaches to Evolutionary Economics*. Ann Arbor: University of Michigan Press.
- YAARI, MENACHIM E. 1987. "The Dual Theory of Choice Under Risk." *Econometrica*, 55: 95-115.

WEB SITES

<http://www.rand.org>
<http://cowles.econ.yale.edu>

<http://www.gmu.edu/departments/economics/facultybios/smith.html>
<http://nobelprize.org/economics/laureates/2002/smith-autobio.html>

SAMENVATTING

Tot voor kort is 'het wetenschappelijke experiment' een onderbelicht onderwerp in de wetenschapsfilosofie. Het is daarom geen verrassing dat een analyse van het experiment waarin processen van kennisproductie zijn uitgewerkt in samenhang met de waarde van experimenteel vergaarde kennis nog ontbreekt. Het doel van deze studie is om een analytisch kader te ontwikkelen dat het interne proces van kennisproductie verklaart en fundamentele epistemologische vragen ten aanzien van 'het wetenschappelijke experiment' beantwoordt. Dit kader wordt dan toegepast op het vakgebied van experimentele economie.

Het in deze studie voorgestelde kader brengt twee aspecten van kennisproductie - aspecten die tot nu toe onvoldoende zijn uitgewerkt in bestaande epistemologische studies naar experimenten - bij elkaar. In deze analyse van de processen waarmee experimentele producten worden geproduceerd en bepaald wordt de rol van zowel "de natuur" als "de sociale wereld" benadrukt. Dit perspectief verklaart de titel van deze studie. De epistemologie van het experiment die in deze studie wordt gepresenteerd is betiteld als een sociale epistemologie om de vaak genegeerde sociale dimensie van het experimentele proces van kennisproductie te benadrukken.

De toepassing van het analytische kader op experimentele economie biedt niet alleen een algemeen overzicht van experimentele economie en de aangrenzende onderzoeksdomeinen, maar ook een accurate beschrijving van specifieke onderzoekprogramma's, experimenten, experimentele technieken en resultaten. Deze diepgaande analyse maakt het mogelijk om de kenmerkende eigenschappen, de legitimerende argumenten en de centrale discussies van de experimentele methode binnen de economische wetenschap te identificeren. Op deze wijze biedt het een volledig en precies beeld van experimentele economie.

De analyse toont ook een spanning die inherent is aan de experimentele methode. De methodologische grondslagen die het mogelijk maken om een micro-economisch experiment in het laboratorium op te bouwen beperken de bijdrage van het belangrijkste epistemologische element van experimentele economie, namelijk de gedragingen van de experimentele subjecten. Echter, deze studie laat zien dat deze spanning kan worden verminderd door de sociale dimensie van kennisproductie. De collectieve aard van experimentele economie bevordert het ontwerpen van nieuwe experimenten waarin de rol van de experimentele subjecten worden vergroot. Deze sociaal-epistemologische analyse van experimentele economie schijnt ook licht op de terugkerende debatten over experimentele economie die – wellicht door het beginstadium van haar ontwikkeling – al te vaak vooroordelen en defensieve reacties bevatten.

CURRICULUM VITAE

Ana Cristina Cordeiro dos Santos was born in Lisbon, Portugal, in 1971. She received her B.Sc. degree in Economics from Technical University of Lisbon, in Portugal, in 1994, and a MA degree in Social Policy from Roskilde University, in Denmark, in 1995. Since 1996 she has been a teaching assistant at Instituto Superior de Ciências do Trabalho e da Empresa (ISCTE), in Lisbon. She obtained a MPhil degree in Philosophy of Economics at the Erasmus Institute for Philosophy of Economics, Erasmus University Rotterdam, in 2001. She completed her Ph.D. in Philosophy of Economics at the same institute.

E-mail: anacsantos@iscte.pt