Never the Single Measure, Essays on Taxes, Careers and Surveying Corporate Managers.

Sander Renes, LL.M. M.Sc.

ISBN: 978 90 361 0403 6

This book is no. **591** of the Tinbergen Institute Research Series, established through cooperation between Thela Thesis and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

Never the Single Measure, Essays on Taxes, Careers and Risky Ventures

Nooit een enkele maatregel, essays over belastingen, carieres en risicovolle beslissingen

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Erasmus Universiteit Rotterdam
op gezag van de rector magnificus

Prof.dr. H.A.P. Pols

en volgens besluit van het College voor Promoties.
De openbare verdediging zal plaatsvinden op

donderdag 27 november 2014 om 9.30 uur

door

Sander Renes
geboren te Amsterdam.

ERASMUS UNIVERSITEIT ROTTERDAM

Promotiecommissie:

Promotor:	Prof.dr. B. Visser
Overige leden:	Prof.dr. B. Jacobs
	Prof.dr. J. Potters
	Prof.dr. O. Swank

Voorwoord

Ergens in mijn derde jaar aan de universiteit, tijdens een vak over economisch beleid, vroeg Prof. Swank, of Otto zoals ik hem nu ken, me of ik wel eens nagedacht had over het halen van een Ph.D. Als ik mijn bachelor thesis goed deed, dan was er waarschijnlijk wel een promotie plek op de ESE te vinden. Mijn doctoraal is nu 5 jaar in de maak, blijkbaar deed ik het goed.

Zoals elk traject op de universiteit wordt de doctoraal studie afgesloten met een geschrift, een dissertatie. Kijkend naar de 4 hoofdstukken die de mijne vormen, kan ik enige gemengde gevoelens niet ontkennen. Een van de belangrijkste redenen om een Ph.D. te doen was het excuus dat het me gaf om op de Uni te blijven en mijn tijd aan de Erasmus zit er nu op. Hoewel ik weer een volgende universiteit gevonden heb, voelt na 9 jaar in Rotterdam de Erasmus als dé Uni. Ook voelen deze 4 hoofdstukken als een incompleet verhaal. Het omvat 'maar' 3 van de projecten waar ik aangewerkt heb, geen aanbestedingen, geen rechten, geen risico in toernooien, of iets van de andere projecten die ik (nog) niet af heb kunnen maken. Bovenal is er een hoeveelheid trots. Dit boekje bevat vier artikelen waar ik mijn naam onder, of zoals gebruikelijk in Academia, boven durf te zetten. Alle vier kunnen ze een bijdrage leveren aan de literatuur en ik hoop dat ze hun plek snel vinden.

Hoewel dit mijn boekje is, heb ik de artikelen niet alleen geschreven. Ik had het geluk dat ik aan de Erasmus veel mensen heb kunnen vinden om mee samen te werken. Niet alleen zijn mijn papers daar beter van geworden, ik had het ook nooit volgehouden als ik 5 jaar lang in mijn uppie op een kantoor had gezeten. Voor hun inzet in het onderzoek en voor de mogelijkheid om zo nu en dan binnen te vallen ben ik Floris, Bauke, Bert, Rene en Philip Hans veel dank verschuldigd.

Als het op binnen vallen aankomt heb je aan Floris sowieso een goede. In de afgelopen jaren zijn we zo vaak bij elkaar het kantoor ingestormd, dat onze kantoorgenoten zich meer dan eens afgevraagd hebben of we niet beter een plek bij konden zetten. Floris is degene die me geïntroduceerd heeft in de wereld van Mirrleesiaanse belasting modellen en hun eigenaardigheden. Deze eigenaardigheden hebben we onderzocht en dit leverde de twee hoofdstukken op die ons beider Job-market paper zouden worden. Na deze twee papers weet ik een ding zeker, als er een probleem is in deze literatuur dat ik niet samen met Floris kan oplossen, dan is het een probleem dat ik niet kan oplossen (al zou het me niet verbazen als het hem later als nog lukt). Dit is een samenwerking waarvan ik hoop dat we hem nog lang doorzetten, ook al zullen we het nu over een paar grenzen heen moeten doen.

Bauke was behalve mijn coauteur ook mijn begeleider. Sinds mijn master scriptie is hij bezig geweest me bij te sturen. Dit deed hij niet door elke week te kijken of ik wel genoeg geschreven had, maar door er te zijn als ik om hulp vroeg. Aangezien ik eigenwijs genoeg ben om dat niet al te vaak te doen, moest hij zo nu en dan vragen waar ik mee bezig was om zeker te maken dat het mijn dissertatie was en niet iets anders. Deze losse begeleidingsstijl is waarschijnlijk niet voor iedereen geschikt, maar ik zou het niet anders willen. Ik heb van weinig dingen meer geleerd dan van mijn pogingen om mijn verschillende onderzoeken aan de gang te krijgen. Bauke gaf me de mogelijkheid mijn eigen weg te vinden en daarvoor ben ik hem zeer dankbaar. Op de ESE en bij het Tinbergen instituut ben ik veel mensen tegengekomen die er voor gezorgd hebben dat onderzoek doen geen eenzame aangelegenheid was. De vakken bij Tinbergen werden een stuk prettiger doordat ik ze samen kon doen met Rogier, Sven en Oke. Mijn kantoorgenoten Nick en Jugo en de laatste vier jaar Wim en Saskia zorgden er voor dat er altijd leven was in het kantoor. Op de gang, in seminars en als ik een probleem wilde bespreken, kon ik altijd rekenen op de input van Barbara, Otto, Bauke, Suzanne, Dana, Josse, Vladimir, Otto, Max, Robert, Aart, Benoit, Jurjen, Bas, Olivier, Rey, Kyle, Zara, Heiner en de vele andere collega's die het H-gebouw zo'n prettige werkplek maken. Ik ga jullie missen.

Probably the most intense part of my Ph.D. training were the four months I spent at UC Berkeley at the invitation of Shachar Kariv. For me this was the first real experience at a university other than the Erasmus, and an amazing experience it was. The campus life in the US is a bit more impressive than around these parts, more seminars, more renowned speakers, more international visitors, and an incredible martial arts program made it a time to remember. I've learned a lot from the professors and students in Berkeley alike, and most of all from John Morgan. He gave me quite the wake-up call about the academic job market, discussed my research with me, and wrote me a reference letter for the job market at the end of it all. Although he would have doubtlessly bested me, I am sorry we never got to play that game of golf.

Een bijzondere groep bij Algemene Economie en Tinbergen vind ik de (administratieve) staf. Zoals zoveel academici ben ik organisatorisch een beetje rommelig. Gelukkig kon ik als ik een deadline had, een computer probleem had, iets raars had met een experimenteel lab, of gewoon hulp nodig had, altijd op de staf rekenen. Ik ging langs de betreffende personen en dan werd er een oplossing gezocht. Na enige aanvaringen met andere, duidelijk minder hulpvaardige, administraties ben ik heel blij met de hulp die ik van jullie gehad heb. Marrianne, Milky, Thea, Jany, Ankimon, Arianne, Judith, Carine, Ine, Caroline, Christina , Sytske, Jan, Urne, Gino en de vele anderen, zonder jullie hulp over de jaren heen had ik het waarschijnlijk een stuk zwaarder gehad en ik had in ieder geval ergens vertraging opgelopen, dank voor alles.

Op het persoonlijke vlak zijn het ook een paar interessante jaren geweest. Zonder de steun van mijn ouders, Alex en Joke, mijn zus, Linda, en mijn vriendin, Jacq, had ik de eindstreep nooit gehaald. Ondanks dat jullie niet altijd volgden waar ik inhoudelijk mee bezig was, waren jullie bereid om mijn monologen aan te horen en waar mogelijk te helpen. Misschien nog wel belangrijker, jullie waren er altijd om me met beide benen op de grond te houden, zodat ik niet verdwaalde in een wereld van multi-dimensionale tweede orde afgeleiden. Halverwege de rit kwam er een extra zorg bij. Na de opkomst van diabetes had ik wel wat tijd nodig te herstellen. Het was erg prettig om te merken dat ik op de mensen om heen kan rekenen. Mijn vrienden van de scouting, RSG, Lambiekske, de SV en mijn collega's (de lijst is iet wat lang om ze een voor een te noemen) hielpen me de draad weer op te pakken. In het bijzonder heb ik veel gehad aan Evey en Foeke die me lieten zien dat er goed mee te leven valt en Frank die me mee nam naar Huk-tti. Op de vloer van Huk-tti, onder Sa Bum Rob Salm en tussen mijn medeleerlingen heb ik een stuk vertrouwen in mezelf terug gevonden. Kwan Jang Nim, Sa Bum nim, Ko dan Ja nim, You Dan Ja nim en You Kup Ja nim, Kamsa hamnida!

For the last step towards becoming a doctor, Bauke and I found Arno Riedl, Bas Jacobs, Jan Potters, Nicola Pavoni, Otto Swank and Victor Maas willing to serve as the defense committee. Their initial remarks have already improved the essays in this dissertation and their questions at my defense will doubtlessly induce more changes. Thank you, o highly learned opponents.

Op de job market had ik niet alleen John Morgan's referentie brief, ook Otto Swank, Bauke Visser en Bas Jacobs waren bereid om hun vertrouwen in mij op papier te zetten. Met deze referenties (en de papers met Floris) op zak ben ik naar de job market gegaan. Ik heb een plek gevonden als Post-doc in Mannheim, een van de beste universiteiten in Europa en bovenal een plek waar ik de komende jaren verder kan leren, want hoewel ik mezelf hopelijk snel doctor mag noemen, verwacht ik voorlopig nog niet uitgeleerd te zijn.

Contents

Voorwoord **v**

1 Introduction **1**

2 When a Price is Enough **5**

2.1 Introduction 5
2.2 Related Literature 11
2.3 The Model 13
2.4 Failure of the Mirrleesian implementation: A Simple Example 21
2.5 Implementation through Taxation 25
2.6 Concluding Remarks 32

3 As Easy as ABC? **35**

3.1 Introduction 35
3.2 Related Literature 40
3.3 The Model 42
3.4 The Second Best Allocation: A First-Order Approach 47
3.5 The ABC Formula 52
3.6 The Atkinson-Stiglitz Theorem 54
3.7 Interdependencies in the Tax System in Mirrleesian and New Dynamic Public Finance 56
3.8 Bunching 59
3.9 Concluding Remarks 60

4 Who Dares? **63**

4.1 Introduction 63
4.2 Related Literature 65
4.3 Methodology 67
4.4 Empirical predictions 70

4.5 Analysis and results . . . 72
4.6 Conclusions . . . 79

5 Career Concerned Managers 83
5.1 Introduction . . . 83
5.2 The Model . . . 87
5.3 Experimental design and hypotheses . . . 91
5.4 Experimental Results . . . 95
5.5 Robustness analysis . . . 106
5.6 Conclusion . . . 110

A Appendices When a Price is Enough 113
A.1 Proof of Lemma 2.3 . . . 113
A.2 Proof to Proposition 2.4 and corollary 2.5 . . . 114
A.3 Proof to Proposition 2.9 . . . 116
A.4 Example . . . 117

B Appendices As Easy as ABC? 121
B.1 Proof of Proposition 3.1 . . . 121
B.2 Proof of Proposition 3.3 . . . 123
B.3 Proof of Proposition 3.7 . . . 124

C Appendices Who dares? 125
C.1 The Questionaire . . . 126
C.2 Summary Statistics . . . 132

D Appendices Career Concerned Managers 133
D.1 Message coding . . . 133
D.2 Experiment Instructions . . . 135

Summary conclusions 141

Nederlandse Samenvatting (Summary in Dutch) 145

Bibliography 149

List of Tables

Chapter 4

4.1 Estimation results of models (1) and (2) 75

4.2 Estimation results of models (1') and (2') with the confidence metric included 78

Chapter 5

5.1 Parameter setting of the experiment 92

5.2 Subject Characteristics 95

5.3 Formation of reputation 97

5.4 Statements as a function of signals 99

5.5 Investment prob. per treatment 99

5.6 Investment prob. as a function of signals received 100

5.7 Committees learning 101

5.8 Percentage of Closed Ranks 102

5.9 Statements made 103

5.10 Signals and Closed fronts 103

5.11 Efficient use of information 106

5.12 Contrasting two committees 107

5.13 Effect of background variables 109

Chapter C

C.1 Summary Statistics 132

Chapter D

D.1 Coding of messages in chat box 134

List of Figures

Chapter 2

2.1 An optimal allocation and a budget constraint. 23
2.2 Utility of a couple faced with the budget constraint in Figure 1. 23

Chapter 4

4.1 self-reported willingness to take risk 73
4.2 Perceived willingness to take risk 74
4.3 The distribution of confidence in the board 77

Chapter 5

5.1 Time line, CM stands for committee member. 88
5.2 Matching schedules 92
5.3 Evolution of Statements and Closed fronts 98
5.4 Distribution of Difference 104

Chapter 1

Introduction

During my time as a Ph.D. student I got interested in phenomena that are most easily characterized as 'gaming' behavior. As a society we have developed extensive systems of rules and institutions to achieve desirable outcomes. These institutions take care of every aspect of our lives, from the interest we get on the bank, to the distribution of crime and punishment and the regulation of radio frequencies for the use of tv, radio and cell-communications. The presence of rules allows us to verify, ex post, whether the correct procedure was followed, and thus ensure that (governmental) power is not abused. They create a system in society that can be used to hold people responsible for their behavior and thus form an essential part of an ordered society.

All these rules and procedures, however, have a flipside; they limit the decision maker's flexibility to respond to individual circumstances and thus make his responses exploitable. By creating predictability in actions, they open up the system to gaming by informed agents. Because of the sheer size of institutions (like governments and regulated markets) in most western societies and the importance of institutions in regulating our lives, gaming can take many forms. In this dissertation, for instance, I look at the possibility of misreporting in the tax-system as well as spurious investment to distort the quality of important information. However, similar stories could be made if one looks at the bribery scandals that plague the FIFA in the current world-cup (not to mention the 2022 one in Quatar), to lobbying of special interests, to the presence of clearly distorted bids in procurement procedures. In all cases some party with an interest and private information exploits vulnerabilities of the decision making process by strategically presenting (or hiding) information, or outright buying influence.

For gaming to exist there has to be a clearly defined, regulated decision process. The process has to be predictable, preferably institutionally fixed, like procedures regulated by law. This will mean that a researcher interested in gaming, will often encounter the law. For someone like me, with a background in both Economics and in Law, the concept of gaming is particularly interesting; it allows me to combine my background knowledge from both fields. Gaming is not a topic on its own, instead it is a problem that can occur in many different settings. This dissertation therefore does not present the reader with a clear singular topic, instead this booklet contains 4 papers that focus on aspects related to gaming in one way or another. The first and second chapter are about the game, they study particular game forms and why these forms are (not) susceptible to gaming. The third and fourth chapter deal with the players, they study empirically how people respond to certain stimuli to see whether or not the assumptions we make about individual behavior are true, or at least close enough to reality to lend our predictions the necessary credence.

The first two papers focus on the tax-system. Most people will realize that tax avoidance or tax evasion can be constructed as examples of gaming. There is one problem with this realization however: it ignores the goal of the tax authority. In a situation with a government that cares for his constituents chapter 2 shows that the constituents, you and me, should not have a reason to duck our responsibility and game the system. If the allocation, the distribution of burdens and benefits, is meant to maximize some concave sum of individual utility, i.e. make us all better off on average, then we all have good reason to cooperate. In effect it shows that for gaming to be a problem two conditions have to be met, the individuals have to disagree sufficiently with the one designing the allocation and they have to have a sufficiently rich set of actions. In short, only if individuals have both reason and opportunity to deviate from the planner's plan, will gaming be a problem. Since in most models studied in the literature individuals have neither reason nor opportunity to deviate, the results show that most solutions found in the literature are implementable through our tax-system. Because gaming is not a problem in these models, the optimal distribution can be implemented through a tax-system that does not use prohibitions or rule-based restrictions, a set of appropriate prices is all that is required of the planner.

Chapter 3 completes the picture that the previous chapter started. It deals with a planner that does not want to create incentives for individuals to game the system, but whose environment is rich enough for individuals to have the opportunity to game his system. The individuals can differ in many aspects and make many different choices and the planner wants to steer them towards an allocation that maximizes total welfare. The differences between people, however, imply that they will not all want to make the same choices. It is exactly this difference between the choices of different individuals that allows the designer to identify different individuals and thus assign them different combinations of goods and bads. In effect, we show that the entire problem can be seen as an informational problem. No matter in how many aspects individuals differ and how many choices they make, as long as the planner wants to design an allocation that identifies individuals' preferences, the allocation will be shaped by the same forces. These driving forces can be conveniently summarized by a multiplication of three factors, known as A, B and C. On such allocations the planner balances the value of information (A) and redistribution (B) gained by marginally distorting a particular choice, with the efficiency cost of this distortion (C). The recovery of this formula allows us to produce generalizations of the famous elasticity approach of Diamond and Saez, a generalization of the Atkinson-Stiglitz theorem and the no-distortion at the top result, while we provide a partial bridge between the complex New Dynamic Public Finance and the much studied classical Mirrlees model.

Chapter 4 focuses on some of the most important decision makers in our economic system, the corporate executives. Especially since the financial crisis of 2008 many discussions in the press and in politics have focused on their role and tasks. A lot of attention has been given to their salaries, most notably to the bonuses they receive and whether or not the use of such large bonuses was warranted by their added value. The discussion also got attention in academia, but here it became a bit single sided in nature. The scientific nature of the discussion sets methodological constraints that limit the discussion. Due to the difficulties of showing their added value scientifically and the difficulties of bringing together and properly questioning the executives, they have hardly been part of the academic discussion. Chapter 4 reports on the results of a survey under Dutch executives. It combines some of the methodology of experimental economics and survey research through an internet based dynamic web-page. The web-page remembers the answers given by individual respondents, which allows the researchers to alter the subsequent questions dynamically, and thus use every bit of information available about the respondents to tailor the questions to the respondent. Our survey first asks about the job of the individual respondents and then uses the responses given to approach the respondents in their professional role. Combined with exogenous variation in the questioned asked, this allows us to show the effects of professional roles on individual decision making, while diminishing the usual confounding factors of endogeneity and observational biases. The results of the survey show that, against our own and our respondents' expectation, in the boardroom it is the CFO that is most willing to take risk, followed by the CEO and (at quite some distance) the Non-Executive or supervisory Directors.

The last chapter deals with a direct test of theory. The theory predicts that experts, informed agents that are hired because they are believed to be capable, will try to manage their reputation. The public wants to hire the best expert since those have the information needed to make the best decision. This implies being perceived capable means an expert gets more jobs in the future and thus is a direct career concern. To manage their reputation these experts need to know what the public expects a capable expert to do. Since the public can use many signals about an expert's ability, e.g. prior decisions, public statements, or the opinion of other experts, this is not a trivial problem. At the other side of the problem we have the public, faced with the task of finding the most capable expert they have to weigh in all the evidence, and figure out what signals to trust and which to ignore. Nash equilibrium requires that while expert tries to make the public believe he is able, the public tries to undo the manipulation of its believes by accounting for the manipulation when weighing the evidence. Since there are so many signals to choose from, this gives both sides a complex task, and it is far from obvious that they the experts will succeed in using the signals that are valued by the public, while it is equally uncertain whether the public will understand what signals to use. In our experiment we allow for a relatively rich information structure, we give half of our subjects the role of experts that produce signals about their ability, while we ask the other half to judge the experts. The theoretical predictions clearly indicate what signals should contain information about ability, and thus should be used by the public and which should be ignored. In our experiments the signals that should contain information according to theory clearly contain more information than the signals that should not. The public responds by placing more value on the informative signals than on the other signals. Although both sides clearly deviate from equilibrium predictions, the theory is accurate enough to help us identify situations in which career concerns can be costly, without having to know what the correct decision is, that is without access to the information only the experts have.

Chapter 2

When a Price is Enough: Implementation in Optimal Tax-Design[1]

2.1 Introduction

Governments in modern welfare states have developed elaborate systems of taxes, benefits, and subsidies (further: tax system). These tax systems are designed to raise revenue for the provision of public goods, to redistribute resources from the fortunate to the less fortunate, to insure constituents against adverse outcomes when insurance markets fail, and to correct for externalities. Tax systems combine non-linear tax rates on labor income, capital income and commodities, with non-linear subsidies on healthcare, housing, education and other observable choices of the constituents.

The formal study of optimal non-linear tax systems was pioneered by Mirrlees (1971, 1976). In his model agents are heterogeneous in their earnings ability. They make several observable choices, such as labor income, portfolio and savings choices, and consumption levels of various goods. The social planner wants to redistribute from agents with high to agents with low earnings ability. However, earnings ability is private information, and

[1]This chapter is based on Renes and Zoutman (2014b). We would like to thank Felix Bierbrauer, Eva Gavrilova, Aart Gerritsen, Yasushi Iwamoto, Bas Jacobs, Laurence Jacquet, Etienne Lehmann, John Morgan, Dominik Sachs, Dirk Schindler, Dana Sisak, Bauke Visser, Casper de Vries and Hendrik Vrijburg for useful suggestions and comments on an earlier version of this paper. Furthermore, this paper benefited from comments and suggestions made by participants at the 2011 Nake Conference, Utrecht, the 2013 CESifo Area Conference on Public Economics, Munich, the 69th IIPF Conference, Taormina; ; and seminar participants at the Erasmus School of Economics, the Norwegian University of Science and Technology, the Norwegian School of Economics, the University of Konstanz and the Centre for European Economic Research. All remaining errors are our own.

hence, the first-best allocation is not attainable. Instead, the main problem in optimal taxation is to find a tax system that implements the social planner's second-best allocation in a market-based economy.

The technique, pioneered by Mirrlees, to identify the optimal tax system is to split the problem into two subproblems. The first subproblem is to characterize the best attainable outcome given the informational restrictions, the second-best allocation, through a direct mechanism. The second subproblem is to identify a tax system under which the second-best allocation is a market equilibrium, and hence, a tax system that implements the second-best allocation in a market-based economy.

There is an extensive literature that studies the first subproblem and describes the properties of the second-best allocation. In addition, Mirrlees (1976) proposes to solve the second subproblem of designing the optimal tax system by equating the marginal tax rates to the optimal wedge between the marginal rates of substitution and transformation (further: equating marginal tax rates to wedges). In his model agents differ only in their earnings ability, such that equating marginal tax rates to wedges yields a separable tax system in which the marginal tax rate on each good only depends on the consumption of that particular good. The Mirrleesian tax system has become the canonical solution in, and has subsequently been applied throughout, the literature on optimal non-linear taxation (see e.g. Atkinson and Stiglitz, 1976, Diamond, 1998, Saez, 2001, Bovenberg and Jacobs, 2005, and Choné and Laroque, 2010).

By equating the marginal tax rate to the wedge, the Mirrleesian tax system solves the first-order conditions for implementation. However, to the best of our knowledge, general conditions under which this tax system satisfies the second-order conditions for market implementation have never been derived. To understand this issue in more detail note that a tax system can only implement the second-best allocation in the market if no agent can increase his utility by deviating from the bundle that would have been assigned to him in the direct mechanism. However, agents typically have more choices available to them in a market economy than they have in a direct mechanism. In the direct mechanism of Mirrlees (1976, sect. 3), for example, agents can choose which earnings ability they report to the planner. Based on this reported ability the planner assigns each agent a bundle of choice variables, such as labor effort, consumption levels and savings. On the market the agent can make each of these choices separately, only constrained by his budget. The combination of these choices allows agents to create bundles that are not available to them in the direct mechanism. Hence, even if an allocation is incentive compatible in the direct mechanism and the tax system satisfies first-order implementability conditions, it does not follow that the tax system implements the allocation in a market economy. The bundles that are not part of the second-best allocation are known in the literature as "joint" or "double" deviations (see e.g. Kocherlakota, 2005). The potential for these joint deviations implies that second-order implementability constraints on the market are generally more stringent than the incentive-compatibility constraints in the direct mechanism.

This issue has not gone entirely unnoticed. Specifically, the principle of taxation derived by Hammond (1979) proves that for any incentive-compatible allocation there exists at least one tax system that satisfies first- and second-order implementability constraints.[2] This tax system stops agents from creating joint deviations in the market by simply prohibiting such deviations. The incentives created by this tax system are mathematically equivalent to the incentives in a direct mechanism. Hence, incentive compatibility in the direct mechanism implies implementability through this tax system.

Hammond's result is of importance to the study of optimal tax systems because it shows that there always exists at least one tax system that can implement the second-best allocation in the market. However, restricting the choices of economic agents to the choices available to them in a direct mechanism effectively removes the benefits of free choice and limited administrative costs associated with a market mechanism. Therefore, this tax system has limited appeal to policy makers in market economies.[3] Moreover, the principle of taxation does not tell us whether less restrictive tax systems, like the Mirrleesian tax system, can implement the second-best allocation.

In addition, in recent literature several papers have focused on identifying the second-best allocation when agents differ in multiple dimensions (see e.g. Cremer et al. (2001), Saez (2002a), Kleven et al. (2009), Choné and Laroque (2010), Renes and Zoutman (2014a), Rothschild and Scheuer (2014) and Jacquet and Lehmann (2014)). However, to our knowledge the second subproblem of designing the optimal non-linear tax system in these models has not yet been addressed.

In this paper we address this gap in our understanding of tax design. We derive conditions under which the Mirrleesian tax system, and indeed any tax system that satisfies first-order implementability conditions, successfully implements the second-best allocation in the market. These results hold independent of the dimension of heterogeneity. In models of optimal taxation where these conditions apply the optimal tax system follows directly from the optimal wedges in the second-best allocation. If the planner adjusts

[2] A similar result has been derived in Maskin (1999) for any countable number of agents.

[3] The name 'principle of taxation' stems from later applications of Hammond's result to taxation by e.g. Rochet (1985), Guesnerie (1995) and Bierbrauer (2009). However, the principle can also apply to the multi-product monopolist pricing problem and auction design (see e.g. Armstrong, 1996). In such a setting the application leads to perfectly realistic implementations, since the monopolist and the auctioneer can choose what (not) to produce/sell and how to bundle their goods.

relative prices in the market by equating taxes to wedges, this suffices and the planner does not need to restrict the choice space of agents in the market.

The analytic starting point of this paper is a set of agents that differ in multiple continuously distributed characteristics, such as earnings ability and tastes. The agents make multiple choices pertaining, for instance, consumption levels of several goods and income from labor and capital. Preferences are described by a completely general utility function. The economy is guided by a planner who wants to implement a second-best allocation of which we know three things: i.) it satisfies the economy's resource constraint, ii.) it is incentive compatible in a direct mechanism, iii.) it maximizes a general welfare function under the constraints given by the other two conditions. It is well-known in the literature that the derivation of the second-best allocation under multi-dimensional heterogeneity of agents is technically very complex. In the companion paper, Renes and Zoutman (2014a), we take up this issue and set the first steps towards fully characterizing the second-best allocation with multi-dimensional heterogeneity of agents. In this paper, however, we take the second-best allocation as given, and use the common elements of such allocations to study sufficient conditions for the implementation of second-best allocations in a market economy.

We first show the relevance of this analysis trough a simple, but instructive example. The example shows that, even when agents only differ in a uni-dimensional characteristic, the canonical Mirrleesian tax system may fail to implement the second best. In our example at least one agent prefers a bundle available to him in the market over the bundle assigned to him in the direct mechanism. The planner can only implement its second best by restricting these joint deviations or taxing them at prohibitive rates. Neither the restrictions nor the prohibitive tax rates can be derived from the second-best allocation. Hence, in this particular example designing the optimal tax system is significantly more complex than deriving the second-best allocation.

We proceed by deriving a lemma which describes general conditions under which a tax system implements the desired allocation using standard micro-economic theory. The first-order implementation constraints require that marginal tax rates are equal to optimal wedges. The second-order implementation constraints require that indifference curves are more convex than budget constraints in all linear combinations of the decision variables. Economists can use these implementability constraints to verify whether a proposed tax system implements the desired allocation. That is, after solving the maximization problem of the planner and formulating the entire tax system, it can be verified whether the tax system satisfies these constraints for the proposed allocation. Unfortunately, most optimal allocations in the literature do not have a closed-form solution. Therefore, verification of implementation can only be performed in the special cases that have been simulated. This verification is useful (and necessary) in such simulations, but it does not provide insights in the general properties of optimal tax systems.

Therefore, our main contribution lies in identifying two classes of optimal tax problems in which a tax system can always implement the allocation, provided that tax rates are equated to wedges. In the identified classes joint deviations are never optimal, and hence, the second-order implementation constraints are automatically satisfied. This implies there is no need to impose restrictions, or levy prohibitive tax rates, on the choice space available to agents in the market, and that implementation does not have to be verified ex post.

The first class of models satisfies the following conditions: i.) the allocation is feasible, incentive compatible and on the Pareto frontier, ii.) there are no externalities iii.) tax rates are equated to wedges, and iv.) the resulting tax system does not have an internal maximum. For this class of problems we show that in any allocation that allows profitable joint deviations, at least one joint deviation exists which increases the utility of the agent and weakly increases tax revenue. Such a deviation entails a Pareto improvement over the original allocation. Hence, the initial allocation could not have been on the Pareto frontier. For optimal taxation, the most important implication of this first class of models is that the second-best allocation of a welfarist planner can always be implemented by equating taxes to wedges, since the second best of a welfarist planner always resides on the Pareto frontier.

A simple corollary to this proposition shows that the Mirrleesian tax system can implement the second-best allocation under uni-dimensional heterogeneity of agents. It follows that a separable tax system can implement the second best in models with uni-dimensional heterogeneity provided the four conditions outlined above are satisfied.

Intuitively, if the planner is welfarist the preferences of the planner and the agents are aligned. The objective function of the planner is, *ceteris paribus*, increasing in the utility of the agent. Hence, budget-neutral joint deviations that increase the agents' utility cannot exist in a market mechanism. It is this property of alignment between the objective of the planner and the agents that allows the planner to implement its second best through a relatively simple tax system.

If the tax function has an internal maximum on the allocation, the proposition fails to apply. Intuitively, from such a maximum a deviation that weakly increases tax revenue does not necessarily exist. However, most tax systems that result from equating taxes to wedges are either monotonic or convex in all choice variables and, as such, they usually do not exhibit internal maxima.

More importantly, the proposition fails for non-welfarist planners or in the presence of externalities. The second-best allocation of a non-welfarist planner is not necessarily on the Pareto frontier (see Kaplow and Shavell, 2001). Similarly, in the case of externalities, even if both the utility of the agent and tax revenue weakly increase, this does not necessarily imply a Pareto improvement, since other agents might be adversely affected through the externality.

The second class of problems where a tax system that equates taxes to wedges can always implement the second-best allocation, is when the second-best allocation is surjective onto the choice space. Intuitively, when each bundle that exists in the market is assigned to at least one type in the direct mechanism, the agents' problem on the market is identical to the agents' problem in the direct mechanism. Therefore, incentive-compatibility and implementability constraints coincide. Since we assumed the original allocation was incentive compatible, it must also be implementable in the market. The prime example of this class of models is the Mirrlees (1971) model where agents only differ in earnings ability and their only choice variables are consumption and labor income.

Our main results hold independently of the number of choice variables, and independently of the number of characteristics in which agents differ. As such, our results have strong implications for the extensively studied models of optimal tax design with unidimensional heterogeneity, as well as for the technically more complex study of optimal taxation under multi-dimensional heterogeneity. Moreover, preferences do not have to satisfy a Spence-Mirrlees or single-crossing condition, since our results remain valid even if there is bunching, such that at least two types receive the same bundle, in the optimal allocation.

The conditions we derive are sufficient conditions, not necessary conditions. There are cases outside of these two classes for which a mechanism that equates taxes to wedges (without additional restrictions) suffices to implement the second-best allocation. Implementation cannot be guaranteed ex ante through our propositions in that case. The two identified classes, however, are of enormous importance since they encompass virtually all non-stochastic models based on Mirrlees (1971, 1976) and the generalization of these models in Renes and Zoutman (2014a) and as a result validate almost all tax systems proposed in the literature.[4]

This paper primarily provides a guide to the relatively understudied second step of optimal tax design. More generally, our results show an imperfect link between direct and indirect mechanisms. A central planner that perfectly observes choices and can price/tax

[4] A notable exception is Jacobs and De Mooij (2011) which extends the Mirrlees model with externalities.

them non-linearly, might still want to rely on quotas or legal prohibitions to reach the second-best allocation. This provides some intuition for the existence of (possibly optimal) complexities in the tax systems in modern welfare states. To prevent abuse of social insurance schemes, a central planner may have to restrict the choices of (potential) beneficiaries, and force them to study, apply for jobs, or enroll in debt counseling for instance.

The rest of this chapter is organized as follows. Section 2 discusses related literature. Section 3 introduces the model. Section 4 contains an example that shows the potential problems in implementation. Section 5 derives our main results and section 6 concludes.

2.2 Related Literature

A large and growing literature has been devoted to deriving second-best allocations under multi-dimensional heterogeneity. Early examples of multi-dimensional screening problems include Mirrlees (1976), Armstrong (1996), Rochet and Choné (1998) and Armstrong and Rochet (1999), and recently there have been several applications in the optimal taxation problem, including Cremer et al. (2001), Saez (2002a), Kleven et al. (2009), Choné and Laroque (2010), Renes and Zoutman (2014a), Rothschild and Scheuer (2014) and Jacquet and Lehmann (2014). The tax systems that implement the second-best allocation in the market are usually left implicit. This paper complements this literature by guiding the design of a tax system that implements the second-best allocations in the market.

In recent literature it has become popular to derive the optimal tax system using a perturbation method (see e.g. Roberts, 2000, Saez, 2001 and Jacquet et al., 2013). With the perturbation method one can directly derive the optimal tax system in the market without resorting to the direct mechanism. However, the perturbation method is mainly applied as a heuristic tool, to find first- but not second-order optimality conditions. Moreover, the perturbation method does not formally show that the derived tax system also achieves the second-best allocation. Hence, in practice, articles that use the perturbation method usually supplement their analysis with a formal proof applying the direct mechanism. Therefore, the insights derived in this paper apply equally to articles that use the perturbation method to identify the optimal tax system.

Stiglitz (1987b) studies implementation in a setting with discrete types. His analysis shows that a discrete distribution of characteristics adds another layer of complexity, since the optimal wedge is determined for each type in the economy, but not for the 'holes' between the types. The insights derived in this paper do not apply directly to the case with discrete types, and we leave this issue for future research.

Our paper also relates to the more general problem of implementation theory, which identifies economic mechanisms that can implement the planner's optimal allocation (see e.g. Hurwicz, 1960, Vickrey, 1961, Hurwicz, 1973, Dasgupta et al., 1979, Myerson, 1979 and Maskin, 1999). In many applications it has been shown that the planner should apply restrictions to the choice space of the agents. In particular, in optimal taxation Guesnerie and Roberts (1984) and Hammond (1979) shows that quantity restrictions on some goods can be welfare improving in the general case, even if the planner sets its linear, respectively, non-linear tax rates optimally. Adams and Yellen (1976) and Armstrong (1996) show that the multi-product monopolist should optimally bundle its products in order to maximize its profits, effectively restricting the combinations of goods available to customers. A similar result has been found in the literature on multi-product auctions, where the auctioneer should optimally bundle its products in order to maximize auction revenue (see e.g. Palfrey, 1983, Chakraborty, 1999, Armstrong, 2000, and Jehiel et al., 2007). Moreover, it has been shown that if several goods are procured or auctioned in a single contract, but the planner is uncertain about the exact quantities required at the time of the procedure, the principal should optimally restrict the range of possible bids. This prevents agents from using skewed bids, bids where the price is too high for one good, and too low for another (see e.g. Athey and Levin, 2001, Ewerhart and Fieseler, 2003, and Renes, 2011). We add to the implementation theory literature by deriving two special cases where the planner can implement its optimum in a market mechanism without imposing any restrictions on the choice space of the agents, provided the planner has access to non-linear taxation/pricing. In particular, we show that when preferences of the planner and the agents are sufficiently aligned, the planner does not need to restrict the choice set of the agents. Hence, the optimal restrictions in the applications above may stem from a misalignment of preferences between the planner and the agents.

The New Dynamic Public Finance has generalized the Mirrlees model to a setting where earnings ability follows a stochastic dynamic process.[5] Kocherlakota (2005) shows that in this setting equating the tax rate to the optimal wedge generally does not implement the second-best allocation, and hence, the planner should restrict the choice set of the agents. Albanesi and Sleet (2006) show that even in the simplest setting, where earnings ability follows an iid process, the planner can implement the optimal allocation only if it supplements income taxes by borrowing constraints. We provide an intuition for this result. As Kocherlakota (2010) a.o. argues, in dynamic stochastic models of taxation, savings by an agent in period t create an "externality" on the labor supply decision of future incarnations of this agent through the wealth effect. Our first proposition does not

[5] See Golosov et al. (2007) and Kocherlakota (2010) for an overview of the literature.

apply when there are externalities, and hence, the planner may need to restrict the savings decisions of agents in order to implement its second-best allocation. Intuitively, the externality causes a misalignment between the preferences of the agents and the planner, and hence, restrictions are required.

2.3 The Model

In this section we lay down the formal structure of our model. First, we define the preferences of the agents in the economy. Second, we define the general properties of a second-best allocation. Third, we define the agents' maximization problem on the market. Finally, we discuss the two methods of implementation that are most commonly applied in the literature: the principle of taxation due to Hammond (1979) and the Mirrleesian tax system due to Mirrlees (1976).[6]

2.3.1 Preferences

The economy is populated by a unit mass of agents characterized by a twice-differentiable utility function:

$$u\left(\mathbf{x}, y, \mathbf{n}\right), \tag{2.1}$$

where $\mathbf{x} \in \mathbf{X} \subseteq \mathcal{R}^k$ is a vector of choice variables, $y \in Y \subseteq \mathcal{R}$ is an untaxed numeraire choice variable, and $\mathbf{n} \in \mathbf{N} \subseteq \mathcal{R}^p$ is the type of an individual. Variables in $\mathbf{x}$ can include e.g. effective labor supply, consumption of (different) commodities, or savings. Choice variables $\mathbf{x}$ and y are observable at the individual level and the social planner can tax all choices in $\mathbf{x}$ non-linearly, but cannot tax y.[7] Throughout the paper we will sometimes refer to the choice variables in $\{\mathbf{x}, y\}$ as goods, even though they can be both inputs and outputs to the production process.

We assume that the untaxed good y is a normal good such that $u_y > 0, u_{yy} \leq 0$ for all values of $\{\mathbf{x}, y, \mathbf{n}\}$. Note that the choice of the numeraire variable has no effect on the optimal allocation, since a tax on y can always be replicated by a uniform tax on all goods

[6] The model description below closely follows that of the companion paper Renes and Zoutman (2014a).

[7] To be able to apply non-linear taxes the planner has to observe each choice variable at the individual level. Hammond (1987) and Guesnerie (1995) show that if anonymous transactions in a good are possible, the planner can only tax that good at a linear rate. Mirrlees (1976, sect. 5), Christiansen (1984), and Boadway and Jacobs (forthcoming) extend the non-linear tax framework to allow for goods that can be traded anonymously, and as such should be taxed linearly. Although, the main intuition of our approach remains valid in this setting, formally extending our model to allow for goods that can be traded anonymously is beyond the scope of this paper.

in $\mathbf{x}$. Therefore, the assumption that y is a normal good is equivalent to the assumption that among all the choice variables available to the agents, there is at least one normal good.

The assumption that y is a normal good directly implies that the utility function is non-satiated everywhere. In addition, it eases interpretation of the results in the remainder of the paper. Since y will act as a numeraire, whether a good is taxed or subsidized can be evaluated by direct comparison y. Moreover, marginal preferences for all goods are neatly summarized by the vector of marginal rates of substitution with respect to the numeraire good:

$$\mathbf{s}\left(\mathbf{x},y,\mathbf{n}\right) \equiv -\frac{u_{\mathbf{x}}\left(\mathbf{x},y,\mathbf{n}\right)}{u_y\left(\mathbf{x},y,\mathbf{n}\right)}.$$

Here element s_i is the negative of the marginal rate of substitution for choice variable x_i with respect to the numeraire y. Therefore, s_i represents the marginal utility loss of receiving an extra unit of x_i, expressed in units of the numeraire variable y. In case of the choice of labor supply, for example, the marginal utility of providing an extra unit of labor is usually assumed negative, it then follows that s_i will be positive in our notation.

Each element in the type vector $\mathbf{n}$ is referred to as a characteristic. Characteristics may include variables such as ability and health status as well as taste parameters. We assume $\mathbf{n}$ follows a multi-dimensional differentiable cumulative distribution function $F\left(\mathbf{n}\right)$, with $F : \mathbf{N} \rightarrow [0, 1]$ and a probability density function $f\left(\mathbf{n}\right)$. Both are defined over the closure of the convex type space $\mathbf{N}$.

The type is private information to each individual and unobservable to the planner. Note that we do not restrict ourselves to static models. Different choices can occur in different periods. However, we do assume that both the type and the mechanism used by the planner are revealed to the individuals before their first choice.

Our utility function allows a very general description of preferences. Note, for example, that the conventional utility representation (e.g. Mirrlees, 1971, Saez, 2001) $\tilde{u}\left(y,l\right)$, where l is labor supply and y is income, is a special case of our utility representation.[8] In addition, we do not require the utility function to be (weakly) separable in any of the goods or characteristics as is usually assumed in the literature. Therefore, agents' preferences for a good may depend on all characteristics of the agent, as well as other consumption choices. Moreover, the utility function does not need to satisfy a single-crossing condition, since our analysis remains valid even if in the optimal allocation types are bunched. Finally, we

[8]To see this, assume that gross income $x_1 = n_1 l$ where n_1 is earnings ability. It can readily be seen that this utility function can be rewritten into our form: $\tilde{u}\left(y,l\right) = \tilde{u}\left(y, \frac{x_1}{n_1}\right) = u\left(x_1, y, n_1\right)$.

do not assume any relationship between the number of characteristics p and the number of goods k other than $k, p \geq 1$.

2.3.2 Incentive Compatibility and Feasibility

The problem of optimal non-linear taxation is to find a tax system that implements the second-best allocation of a social planner in a market economy. Following the methodology of Mirrlees (1971, 1976), this requires us to solve two subproblems. The first subproblem is to find the second-best allocation using a direct mechanism. The second subproblem is to find the optimal tax system that implements the second-best allocation in the market. This paper focuses on the second problem, but before we can find the optimal tax system, we first need to outline the properties of a second-best allocation.

A second-best allocation that has been derived through a direct mechanism must satisfy the following three criteria: i.) the allocation is resource feasible, ii.) the allocation is incentive compatible, and iii.) the allocation maximizes the central planner's objective function under the constraints given by the first two conditions.

Formally, these properties can be defined as follows. Let the second-best allocation of goods be denoted by:

$$\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\} \quad \forall \quad \mathbf{n} \in \mathbf{N}.$$

Here $\mathbf{x}^*$ and y^* are functions mapping from the type space to the good space, $\mathbf{x}^* : \mathbf{N} \to \mathbf{X}$ and $y^* : \mathbf{N} \to Y$. We assume $\mathbf{x}^*(\cdot)$ and $y^*(\cdot)$ are both twice differentiable in all their arguments.[9]

Let $\mathbf{X}^*$ denote the image or range of function $\mathbf{x}^*$, and Y^* the image of y^*. We denote by $\{\mathbf{X}, Y\}^*$ the image of the allocation, and its complement by $\{\mathbf{X}, Y\}^{C*}$. This image contains the collection of all bundles assigned to agents in the economy. By definition, the set of assigned bundles are a subset of the goods space, $\mathbf{X}^* \subseteq \mathbf{X}$, $Y^* \subseteq Y$. The sets $\mathbf{X}^*$ and $\mathbf{X}$ (Y^* and Y), are equal if each possible value of choice variables is assigned to at least one agent. Similarly, $\{\mathbf{X}, Y\}^* = \{\mathbf{X}, Y\}$ if all possible combinations of goods are assigned to at least one agent.

Since the economy should be able to produce all goods consumed, we know the optimal allocation must satisfy the economy's resource constraint. We assume the economy's

[9]This assumption makes the rest of the analysis easier. However, by non-satiation of the utility function we know that individual budget constraints are binding, while theorem 1 of Clausen and Strub (2013) guarantees that the FOC binds with equality on interior choices. The combination of these two constraints imply that on any region of the type space that separates our tax system is well defined and our proof structure holds.

resource constraint takes the form:

$$\int_{\mathbf{N}} y^* (\mathbf{n}) \, dF(\mathbf{n}) + R = \int_{\mathbf{N}} q(\mathbf{x}^* (\mathbf{n})) dF(\mathbf{n}). \tag{2.2}$$

In this equation, R is the exogenous revenue requirement of the planner measured in units of y and $q : \mathbf{X} \to Y$ is a reduced form production function that describes the production of y as a function of the total consumption of goods in $\mathbf{x}$. The equation states that total production of y should equal the sum of private consumption and the revenue requirement of the numeraire good. Since the price of the numeraire good is normalized to one, the derivative q_{x_i} represents the individual (negative) marginal rate of transformation between good x_i and the numeraire good y. q_{x_i} Is positive when good x_i is an input and negative when x_i is an output of the production process. We assume weakly decreasing returns to scale such that all $q_{x_i x_i}$ are non-positive. Note that the conventional assumption in the literature that $q\,(.)$ is linear[10] is a special case of our model. The generalization of the production function allows researchers to investigate the case where individuals face decreasing returns to scale because they, for example, have a lower marginal productivity as they supply more hours of labor. An allocation that satisfies condition (2.2) is said to be feasible.

Incentive compatibility requires that in a direct revelation game each agent reveals his type. In a revelation game the planner presents the agents with a menu of choices over bundles, $\{\mathbf{x}^* (\mathbf{m}), y^* (\mathbf{m})\}$, from which each agent chooses his preferred bundle by sending the corresponding p-dimensional message about his type, $\mathbf{m} \in \mathbf{N}$, to the planner. An allocation $\{\mathbf{x}^* (\mathbf{m}), y^* (\mathbf{m})\}$ is incentive compatible if each individual truthfully reveals all his unobserved characteristics through message $\mathbf{m}$ and receives the bundle designed for him. That is, incentive compatibility implies the agent maximizes his utility by sending the message $\mathbf{m} = \mathbf{n}$.

Hence, an incentive-compatible and feasible allocation can be defined as follows:

Definition 2.1. *An allocation* $\{\mathbf{x} = \mathbf{x}^* (\mathbf{n}), y = y^* (\mathbf{n})\} \quad \forall \quad \mathbf{n} \in \mathbf{N}$ *is incentive compatible and feasible if each agent truthfully reveals his entire type in a direct mechanism:*

$$\mathbf{n} = \arg\max_{\mathbf{m}} u \left(\mathbf{x}^* (\mathbf{m}), y^* (\mathbf{m}), \mathbf{n}\right) \quad \forall \mathbf{n} \in \mathbf{N}. \tag{2.3}$$

and in addition satisfies constraint (2.2).

[10]See e.g. Mirrlees (1971), Mirrlees (1976), Diamond (1998), Saez (2001) and Jacquet et al. (2013).

Note that incentive compatibility does not imply that there is no bunching of types. Two different types may well receive exactly the same bundle of goods, provided it is in both types best interest to choose this bundle over all other bundles in $\{\mathbf{X}, Y\}^*$.

A second-best allocation must maximize the objective function of the planner among the set of incentive-compatible and feasible allocations defined in definition 2.1. However, at this point we deliberately leave the objective function of the social planner entirely generic, such that the second-best allocation of the social planner may reside anywhere within the incentive-compatible and feasible set. This allows us to study implementation in the most generic setting possible. We will later show, that whether or not an allocation can be implemented, can crucially depend on the assumptions made about the objective function of the planner.

Characterizing the optimal allocation when agents differ in multiple dimensions may be very complex and is beyond the scope of this paper. The common approach is to assume that the second-order incentive constraints are not binding. Mirrlees (1976); McAfee and McMillan (1988); Renes and Zoutman (2014a) describe characteristics of optimal incentive-compatible allocations in the context of a first-order approach where second-order conditions of maximization problem have to be checked after the entire allocation is derived, while Rochet and Choné (1998) deals with the problems caused by second-order incentive constraints. However, in this paper we are not interested in deriving the second-best allocation. We start instead from an incentive-compatible and feasible allocation that satisfies definition 2.1. Following proposition 2 in Renes and Zoutman (2014a) we know that such an allocation can be described by the optimal wedges:

$$\mathcal{W}_i(\mathbf{n}) = q_{x_i}(\mathbf{x}^*(\mathbf{n})) - s_i(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}) \quad \forall \mathbf{n} \in \mathbf{N}. \tag{2.4}$$

These wedges represent the difference between the marginal rate of substitution, s_i, and the marginal rate of transformation, q_{x_i}, for each agent and each good on the second-best allocation. If the wedge is positive the marginal rate of transformation is larger than the marginal rate of substitution such that the distorted consumption level of the good is below its laissez-faire value. If the wedge is negative the consumption of a good is above its laissez-faire value.

2.3.3 Market Implementation

The social planner aims to implement the second-best allocation in a market economy through a tax system. In the market, agents maximize their utility function (2.1) with respect to their choice variables $\mathbf{x}$ and y subject to their budget constraint:

$$y \leq q(\mathbf{x}) - T(\mathbf{x}), \tag{2.5}$$

where the tax system, T, maps from the good space to the numeraire, $T : \mathbf{X} \to Y$. How much a consumer can spend on y depends on his choice of $\mathbf{x}$, the production function $q(\cdot)$ and the tax system $T(\cdot)$.

A tax system implements an allocation if each agent weakly prefers his bundle over all other combinations of goods available to him in the market. This concept is formally defined in definition 2.2:

Definition 2.2. *A tax system $T(\mathbf{x})$ implements an allocation $\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\}$ in the market if each agent selects the same bundle on the market as is assigned to him in the allocation:*

$$\begin{aligned} \{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\} &= \arg\max_{\mathbf{x},y} \{u(\mathbf{x}, y, \mathbf{n}) : y = q(\mathbf{x}) - T(\mathbf{x}), \mathbf{x} \in \mathbf{X}\} \\ \forall \quad \mathbf{n} \in \mathbf{N} & \end{aligned} \tag{2.6}$$

Note that we do not need to check whether the planner's budget constraint is satisfied as long as the tax system is successful in implementing the second-best allocation. By definition 2.1 the second-best allocation is feasible and the non-satiation of the utility function implies that on the market the agents' budget constraints must hold with equality. By Walras' law the planner's budget constraint must, therefore, also be satisfied.

The difficulty of implementability can be understood by comparing definitions 2.1 and 2.2. In the direct mechanism each agent maximizes his utility by sending the optimal p-dimensional message, containing all p characteristics of his type, to the planner. In doing so, he can choose his most preferred bundle in the set $\{\mathbf{X}, Y\}^*$. Agents can deviate by mimicking another type, but they cannot receive a bundle that has not been assigned to any type. However, in the market each agent can choose his optimal bundle out of all points in the choice space $\{\mathbf{X}, Y\}$ within his budget set. The market allows the agents to create new bundles that were not assigned to any type in the direct mechanism. Such a strategy is called a joint deviation since, in order to create a new bundle that satisfies the budget constraint, an agent has to deviate in at least two goods relative to his bundle in the second-best allocation. If a tax system allows for profitable joint deviations, it cannot satisfy definition 2.2, even if the allocation satisfies definition 2.1. Because the market allows agents to create joint deviations, the conditions for implementability are generally stricter than the conditions for incentive compatibility. A graphical example based on the canonical Mirrleesian tax system will illustrate the issue in more detail in the next section. Before discussing the example it is convenient to discuss the most important approaches to implementation in the literature in our notation.

2.3.4 The principle of taxation and the Mirrleesian implementation

We are not the first to study tax systems that meet definition 2.2. The two most prominent approaches to implementation in the literature are the principle of taxation and the Mirrleesian tax system. The principle of taxation, derived in Hammond (1979), shows that for any incentive-compatible allocation at least one tax system exists that can implement it. This tax system has two properties. First, if the agent chooses the bundle $\mathbf{x}^*(\mathbf{n})$ designed for him, he will receive the corresponding value of $y^*(\mathbf{n})$. That is, the tax function satisfies:

$$T\left(\mathbf{x}^*(\mathbf{n})\right) = q\left(\mathbf{x}^*(\mathbf{n})\right) - y^*(\mathbf{n}) \qquad \forall \quad \mathbf{n} \in \mathbf{N}.$$

Second, it restricts agents to making a choice within $\{\mathbf{X}, Y\}^*$. The implementation creates individual budget sets that are restricted to lie within the image of the allocation. This restriction effectively prohibits all joint deviations. This immediately implies that the underlying problems of the agents in definitions 2.1 and 2.2 are isomorph, and hence, the outcome is identical. It follows that the conditions for incentive compatibility and market implementability coincide. Hence, the principle of taxation shows that by sufficiently restricting the choice space available to agents in the market, a tax system can always implement the second-best allocation.

Mirrlees (1976) derives the second-best allocation with multiple goods under unidimensional heterogeneity in earnings ability of the agents. The outcome of the planner's maximization defines the wedges, $\mathcal{W}_i(n)$, for each level of earnings ability, n, and each good x_i. Under certain regularity conditions the consumption level of each good can be used to infer n.[11] That is, there exists an inverse function of the allocation $\mathbf{x}^*(n)$, such that $(\mathbf{x}^*)^{-1}(x_i) = n$ for each good x_i.

Mirrlees (1976) proposes to implement the second-best allocation through a tax system that has the following properties. First, like before, each agent should be able to afford his bundle:

$$T\left(\mathbf{x}^*(n)\right) = q\left(\mathbf{x}^*(n)\right) - y^*(n) \qquad \forall \quad n \in N,$$

[11] To be precise this requires non-satiation of the utility function, the Spence-Mirrlees condition on preferences and a monotonicity constraint on the allocation, as well as single-dimensional heterogeneity.

However, unlike in the principle of taxation, Mirrlees does not limit the choice set to $\{\mathbf{X}, Y\}^*$. Instead, he proposes to set the marginal tax on each good x_i equal to the wedge:

$$T_i'(x_i) = \mathcal{W}_i(n) = \mathcal{W}_i((\mathbf{x}_i^*)^{-1}(x_i))$$

In the literature this approach is known as equating taxes to wedges (see e.g. Kocherlakota, 2005). The tax system is separable in the sense that the tax rate on good x_i only depends on the consumption of good x_i. Additionally, if all $x_i \in \mathbf{X}$ are awarded to a type in the direct mechanism, this defines the unique separable tax system because a marginal tax rate is assigned to each level of the choice variables. Furthermore, the tax system is entirely defined by the optimal wedges in the direct mechanism. Therefore, the optimal allocation in the direct mechanism contains all the information necessary to design the optimal tax system.

Note that the tax system described does not prohibit choices of the agent in any way, apart from setting a budget constraint. Unlike the tax system derived through the principle of taxation, the Mirrleesian tax system does not change the agents' choice set compared to laissez faire, but only affects the relative prices through the marginal tax rates. As a result, the agents can choose joint deviations if they so desire. Therefore, unlike the tax system prescribed by the principle of taxation, there is no easy proof that shows that this tax system can implement all incentive-compatible allocations. Conditions under which this tax system implements the second best are derived in our corollary 2.8 further on.

Unfortunately, Mirrlees (1976) only defined implementation under uni-dimensional heterogeneity. When agents are heterogeneous in multiple characteristics it is generally not possible to create a separable tax system, since the inverse allocation function $(\mathbf{x}^*)^{-1}(\cdot) = \mathbf{n}$ can no longer be written as a function of a single good (see also the discussion in Renes and Zoutman, 2014a). Since the tax system is no longer separable, there may be multiple natural ways to extend the marginal tax rates to consumption bundles that are not awarded to any type in the direct mechanism. Moreover, there might be multiple tax systems that equate the marginal tax rates to the optimal wedges, since the inverse allocation function is generally not unique. Because of this potential multiplicity we do not focus on a particular tax system. Instead, we derive general conditions under which tax systems that affect the relative prices, but do not restrict the choice space can implement the second-best allocation, even if agents are heterogeneous in multiple dimensions. However, given the focus on uni-dimensional heterogeneity of agents in the literature, the most important example is the Mirrleesian tax system.

2.4 Failure of the Mirrleesian implementation: A Simple Example

This section will use a graphical example to show intuitively why the standard Mirrleesian implementation may fail to implement an incentive-compatible allocation. The purpose of the example, depicted in figures 2.1 and 2.2, is to illustrate the issue of implementation in the simplest possible setting. For that purpose the example is highly stylized and very much simplified. The idea is specifically not to give a realistic example, but to show that even in very simple settings, with agents that differ only in their earnings ability, the Mirrleesian tax system can sometimes fail.[12]

The agents in the example are couples that maximize a joint utility function. For simplicity we assume that spouses within a couple have matched perfectly assortative, such that within each couple the spouses have exactly the same earnings ability. Moreover, we assume couples differ only in their earnings ability, such that this is a model with unidimensional heterogeneity of agents, $p = 1$. The spouses have to decide how much time either partner works and how much each of them tends to the household and children. The optimal allocation specifies how much labor income is generated by each spouse, $x_1^*(n)$ and $x_2^*(n)$, and how much each couple consumes, $y^*(n)$, as a function of the earnings ability of the couple. We assume the standard production function $x_i = nl_i$, where l_i represents the labor effort of spouse i. Clearly, high ability couples have to provide less labor to reach a certain income level than low ability couples. Moreover, we assume that couples' marginal utility of consumption is decreasing in consumption, such that the planner has an incentive to redistribute from high to low ability couples.

Since there is only one hidden characteristic for each couple, the bundles assigned to the types in a second-best allocation form a line in $X_1 \times X_2 \times Y$ space. This line is represented by the thick black line in figures 2.1 and 2.2. In figure 2.1 the line is sloping upward, indicating that the planner wants couples with higher ability to earn more in equilibrium. In addition, an increase in gross income, $x_1 + x_2$, leads to a less than one-to-one increase in consumption, y. This shows that the planner is redistributing from higher to lower ability couples. Finally, on any point of the line $x_1 = x_2$, indicating that the planner wants the spouses in each couple to supply the same amount of labor effort. In both pictures the dot represents the bundle assigned to one particular couple in the second-best allocation.

[12] In the appendix we describe one particular set of preferences for planner and the agents for which this allocation is second best.

The hyperplanes show the budget constraint that results from the Mirrleesian implementation. The budget constraint is uniquely determined by equating taxes to wedges on the allocation and assuming it is separable in each choice variable. Each point on the surface represents a bundle of labor income of the husband, x_1, labor income of the wife, x_2, and the corresponding amount of consumption after taxes, y, given by the budget constraint.

In figure 2.2 the vertical axis shows the utility level of the couple whose second-best bundle is located at the dot. The surface represents the reduced-form utility function after eliminating y via the budget constraint, for all combinations $\{x_1, x_2, y\}$ given the tax system imposed. In figure 2.2 we can see that the assigned bundle (dot) marks the highest utility level on the allocation (line), such that the couple prefers their bundle over any of the other bundles in the allocation. The allocation is therefore incentive compatible for this couple, and indeed for all couples, in the direct mechanism. Mimicking another couple, and receiving the corresponding bundle, would decrease their utility level. Moreover, the depicted allocation is feasible.

However, in the market the couple is also allowed to trade working hours of the husband for working hours of the wife and vice versa. In figure 2.2 we can see that such joint deviations give the couple located at the dot more utility than their assigned bundle. In the market this couple prefers bundles that are not part of the planner's second-best allocation over the bundles in the allocation. Therefore, the Mirrleesian tax schedule fails to implement the second-best allocation in this case.

In this particular example, the failure results from divergence in preferences between the planner and agents. Our fictitious couple prefers partial specialization, where one of the partners earns most of the income and the other partner stays at home. This could be rationalized in a model where spouses have increasing returns to scale or increasing utility to specialization. However, in the planner's optimal allocation both spouses work the same number of hours. Such misalignments could have several causes. First, the planner could be paternalistic, forcing households to equally divide the task between both partners, simply because he thinks this is fair. Alternatively, the couples in the economy might have children. The children in turn may be better off with attention of both spouses than with attention of only the stay-at-home spouse. If the utility of the child is not sufficiently weighted in the utility function of the couple, the planner may want to correct the externality and make both spouses stay at home part of the time. Note that both of these arguments can lead to exactly the same first-order optimality conditions for the second-best allocation. This example indicates that both non-welfarist motives and correction of externalities may lead to a failure of the Mirrleesian tax system,

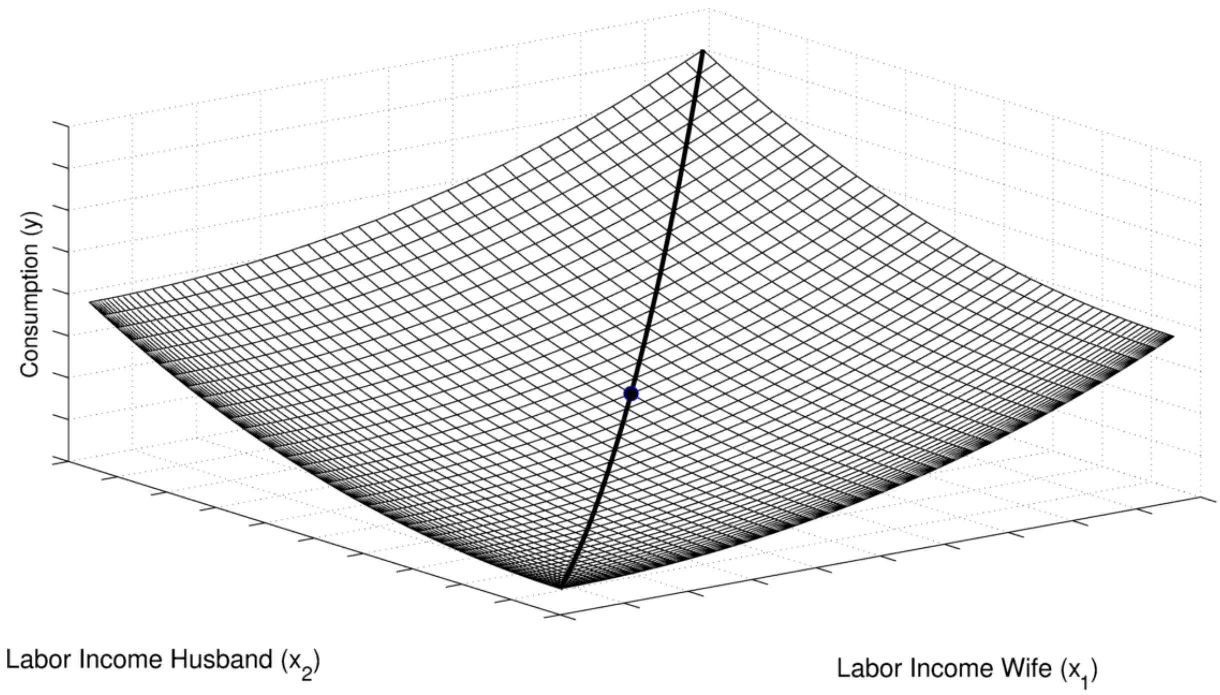

Figure 2.1: An optimal allocation and a budget constraint.

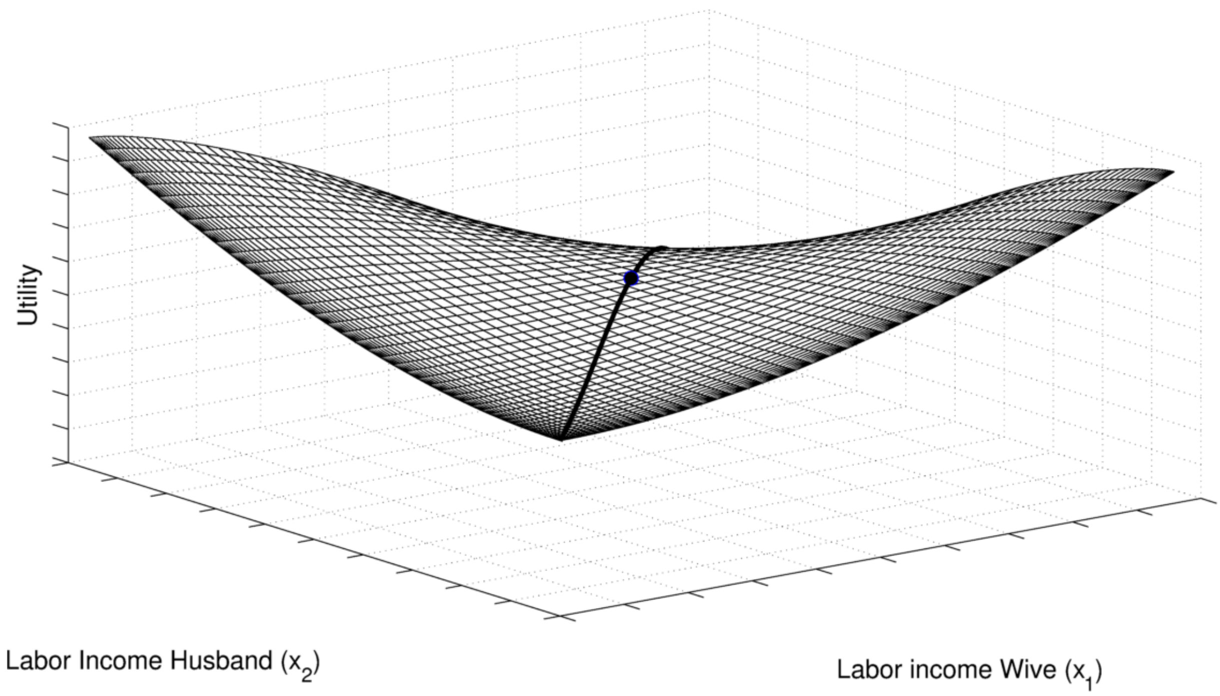

Figure 2.2: Utility of a couple faced with the budget constraint in Figure 1.

independently of the preferences of the agents. In section 5 we prove that absent this type of misalignment any tax system that equates taxes to wedges and does not impose any other restrictions on the choice set of agents, like the Mirrleesian tax system, suffices to implement the allocation.

To complete the picture, note that the solution offered by the principle of taxation is to disallow all consumption bundles that are not on the thick black line. That is, it forces both spouses to work an equal number of hours. In that case the observation that the assigned bundle (the dot) represents the highest utility level on the allocation (line) directly implies implementability.

However, in this example, as in many other cases, implementing the allocation through the principle of taxation may be undesirable. First, the planner should have the practical means to disallow joint deviations. The planner could, for instance, tax joint deviations at an infinite rate, or prohibit joint deviations explicitly, giving severe punishment to trespassers. Such measures, if they are feasible to begin with, are sensitive to small mistakes or a trembling hand. In our example, a slight deviation where one partner works more than the other for some unforeseen reason would be outright prohibited or subject to a severe punishment.

More generally, the planner could tax joint deviations by prohibitive, but non-infinite tax rates. However the calculations required to determine the second-best allocation do not provide us with information about when tax rates become prohibitive. Obviously, this crucially depends on the preferences of the agents. Hence, another set of complicated calculations is necessary to determine whether or not specific tax rates indeed lead agents to refrain from joint deviations in equilibrium.

Additionally, it may be costly and complex to administer prohibitions or prohibitive tax rates. In order to determine whether a specific bundle entails a joint deviation the entire vector of $\mathbf{x}$ has to be known and compared to the bundles in $\{\mathbf{X}, Y\}^*$ and this may be a very costly process, in terms of computational complexity, monitoring and administration.

This example clearly shows that a Mirrleesian tax system cannot implement all second-best allocations, and that the principle of taxation provides policy makers with tax systems that are unrealistic in market economies. Hence, in order to guide tax design for policy makers, we have to find conditions under which more realistic tax systems can implement a given allocation.

2.5 Implementation through Taxation

Now that we have shown that the canonical Mirrleesian tax system may sometimes fail to implement the second-best allocation. We will proceed by deriving sufficient conditions for a tax system to implement the second-best allocation. These conditions are a useful test to verify whether a specific tax system implements a specific allocation. However, in order to perform this test, one first needs to derive both the allocation and the tax system. In many cases an explicit closed-form solution for the second-best allocation does not exist. Numerical solutions are available, but these describe only special cases by definition. Implementation can only be verified for the specific tax system and parameterization studied. Therefore, these explicit solutions cannot be used to say anything about tax systems in general.

To overcome this issue we use to lemma derived in subsection 2.5.1 to derive two classes of problems for which a tax system that equates taxes to wedges always implements the second-best allocation in subsections 2.5.2 and 2.5.3. In the final subsection we discuss the consequence of this result for the existing literature on optimal taxation.

2.5.1 Conditions for Implementation

The sufficient conditions for implementation can be found by formally solving the problem in definition 2.2 in terms of the first- and second-order conditions of optimality of an interior maximum. Although this approach is standard, the non-linear nature of prices in our setting requires some attention. Lemma 2.3 describes conditions for a global, interior maximum:

Lemma 2.3. *An incentive-compatible and feasible allocation* $\{\mathbf{x}^*, y^*\}$ *can be implemented through a twice differentiable tax system* $T(\mathbf{x})$ *if a.e.:*

i.)

$$y^*(\mathbf{n}) = q(\mathbf{x}^*(\mathbf{n})) - T(\mathbf{x}^*(\mathbf{n})), \tag{2.7}$$

ii.)

$$T_i'(\mathbf{x}^*(\mathbf{n})) = \mathcal{W}_i(\mathbf{n}), \tag{2.8}$$

iii.)

$$-\frac{\partial \mathbf{s}(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n})}{\partial \mathbf{x}^*(n)} + q''(\mathbf{x}^*(\mathbf{n})) - T''(\mathbf{x}^*(\mathbf{n})) \lessdot 0. \tag{2.9}$$

where $\lessdot 0$ *denotes that the matrix on the left-hand side should be negative semi-definite.*

iv.)The maximum described in local terms by (2.7) - (2.9) is the global maximum of utility under the budget constraint.

Proof. The proof can be found in the appendix. □

Equation (2.7) ensures that the amount of taxes paid for any bundle of $\mathbf{x}^*(\mathbf{n})$ within the allocation is uniquely determined. If the total tax level $T(\mathbf{x}^*(\mathbf{n}))$ is too high, the tax schedule cannot implement the allocation because people receive too little $y^*(\mathbf{n})$, and vice versa. Equation (2.8) is the first-order condition for a market implementation. It states that marginal taxes are equated to the wedges. There are always as many marginal tax rates in T' as there are goods in $\mathbf{X}$, for all $\mathbf{n} \in \mathbf{N}$. Such that there is always a unique vector of marginal tax rates $T'(\mathbf{x}^*(\mathbf{n}))$ that satisfies (2.8) within any possible incentive-compatible allocation.

In effect, this means that the first order conditions of this problem can always be met and that the solution is unique on the allocation, but undefined for consumption bundles outside of $\mathbf{X}^*$. In our example in figures 1 and 2, this translates to a tax system that is well defined on the line, but undefined everywhere else. Hence, although equation (2.8) tells us how to design the tax system on the allocation, it does not tell us how to extend it to bundles that are not assigned in the second-best allocation. Therefore, in order to design a complete tax system one has to make additional assumptions. In the case of the Mirrleesian tax system this additional assumption is that marginal tax rates are separable in each choice variable. One can think of other solutions, and the optimal tax system is generally not unique. However, any tax system that successfully implements the second-best allocation must also satisfy the second-order implementability constraints, equation (2.9).

Equation (2.9) states that the indifference curves of any linear combination of $\mathbf{x}$'s with respect to y should be more convex than the budget constraint for the same linear combination of $\mathbf{x}$'s. This condition is different from the standard second-order condition of utility maximization with two goods (see e.g. Mas-Collel et al., 1995) in two ways. First, in standard micro-economic theory the budget constraint is linear and hence the condition can be simplified to convexity of the indifference curves. Second, since there are multiple choices, sufficiency requires that the indifference curve of all linear combinations of $\mathbf{x}$ with respect to y are more convex than the budget constraint, with two goods only one such combination exists.

Condition iv.) states that the maximum in the allocation, defined in local terms by equation (2.7)- (2.9), is the global maximum in individuals' utility, for which only sufficient conditions can be given.

2.5.2 Pareto Efficiency and Welfarist Planner

As we have seen in figure 2.2 a Mirrleesian tax system may sometimes allow agents to create profitable joint deviations. In the next proposition we show that this issue does not occur if an allocation is i.) incentive compatible and feasible as in definition 2.1 and, ii.) within this set, on the Pareto frontier, such that no agent can be made better off without reducing the utility of any other agent. The most important implication of this proposition is presented in corollary 2.5, the second-best allocation of a welfarist planner can always be implemented by equating taxes to wedges.

Proposition 2.4. *If an allocation satisfies the constraints given by definition 2.1 and, within this set, resides on the Pareto frontier, then it can be implemented by any tax system,* $T(\mathbf{x})$*, that; i.) equates taxes to wedges on the allocation,* $T_i'(\mathbf{x}^*(\mathbf{n})) = \mathcal{W}_i(\mathbf{n})$*, ii.) does not contain an internal maximum in any vector subspace of* $\mathbf{X}^*$*, as long as condition iv.) of lemma 2.3 is met.*

Corollary 2.5. *The second-best allocation of a planner that maximizes a welfare function of the form* $\int_{\mathbf{N}} W(u(\mathbf{n}), \mathbf{n})\, d\mathbf{n}$ *with* $W_u(\cdot) > 0$ *for all* $u, \mathbf{n} \in \mathbf{N}$*, commonly known as a welfarist planner, can be implemented under the conditions in proposition 2.4.*

Proof. The proof can be found in the appendix. □

Intuitively, consider again our example in figure 2.2. The increase in the reduced-form utility of the couple is approximately the same whether the couple turns “left ” or “right”of the bundle assigned to them. That is, in the example it does not matter which spouse specializes in the official labor market, and which spouse specializes in house work. If tax payments from the couple to the planner are not maximized at their allocated bundle, then one of the two opposite deviation strategies must weakly increase tax revenue as well. Hence, one of the two deviations must be resource feasible. By assumption the utility of other couples is not directly affected by either deviation. Hence, there exists at least one joint deviation that has the following properties i.) it increases the utility of one agent, ii.) it occurs through voluntary transactions on the market and is hence incentive compatible, iii.) it weakly increases the resources in the economy, and iv.) it leaves the utility of other agents unaffected. This deviation must be a Pareto improvement. However, an incentive-compatible Pareto improvement does not exist for allocations located on the Pareto frontier, and hence, we arrive at a contradiction. The proof to the proposition shows formally that this logic is not partial to our example. If taxes are equated to wedges but second-order implementability constraints are not satisfied, then there always exists at least one incentive-compatible deviation that increases the utility of an agent and weakly increases the economy's resources. In our setting such a deviation is always a Pareto improvement.

Proposition 2.4 is important for public economists and policy makers. It shows that if the conditions of proposition 2.4 are fulfilled, the design of the optimal tax system is relatively straightforward. If a second-best allocation is Pareto efficient, the optimal tax system can be designed by equating the marginal tax rates to the optimal wedges. By proposition 2.4 this tax system must satisfy first and second-order implementability constraints, and can hence implement the planner's optimal allocation.

In addition, by corollary 2.5 such a tax system can also implement the second-best allocation of a welfarist planner. This follows from the fact that the second-best allocation of a welfarist planner always resides on the Pareto frontier, as was already shown in Werning (2007) and Brendon (2013). Intuitively, if the planner is welfarist, preferences of the agents and the planner are aligned in the sense that welfare strictly increases in the utility of all agents. Therefore, an allocation which is not on the Pareto frontier cannot be optimal for a welfarist planner.

The conditions of proposition 2.4 are only satisfied if the tax system does not have an internal maximum in any vector subspace of $\mathbf{X}^*$, there are no externalities, the allocation is on the Pareto frontier and the optimum is a global maximum. In practice, the first restriction is almost always satisfied. The sign of the marginal tax rate is equal to the sign of the wedge. In most models of optimal taxation the optimal wedge does not change sign, such that the resulting tax system is monotonic, and hence, the tax system does not have an internal maximum. Even in models where the optimal wedge does change sign, such as Saez (2002b) and Choné and Laroque (2010), it changes sign from negative to positive. As such, the resulting tax system has an internal minimum, but not an internal maximum. We are not aware of any articles in the literature where this assumption is violated.

More importantly, the logic of our proof does not carry over to a model with externalities or to situations where the planner is non-welfarist. With externalities the deviation of any agent can influence the utility of other agents through the externality. Therefore, if a deviation increases the utility of a single agent, and weakly increases tax revenue, this no longer implies the deviation is a Pareto improvement. For instance, if a deviation increases the amount of gasoline purchased by an agent, this may lead to an increase in the utility of the agent as well as the tax revenue of the planner. However, the externality generated by the extra carbon emissions may harm the utility of all other agents in society. In fact, Kaplow and Shavell (2001) show that for a planner with a non-welfarist objective function there exist at least one Pareto improvement for the agents that decreases the objective function of the planner. Hence, even if a Pareto improvement exists, it is unclear whether this increases the objective function of the non-welfarist planner.

The last restriction for implementation is the hardest one to deal with. In the literature only sufficient conditions for a global maximum are known, but several of these classes deserve some attention because of their relevance for the literature.

Remark 2.6. *Condition iv.) of lemma 2.3 is automatically met if preferences are convex in* $\{\mathbf{x}, y\}$*, the production function satisfies Constant Returns To Scale (CRTS) and taxes are (weakly) convex.*

If the tax schedule is weakly convex and the production satisfies CRTS the individuals budget set is convex. Standard results in consumer theory, see for instance proposition 3.d.iii in Mas-Collel et al. (1995), guarantee that the demand correspondence of the indivdual is single valued and this is the unique stationary point, the maximum described must thus be unique.

Remark 2.7. *Condition iv.) of lemma 2.3 is met if condition 2.9 holds for all values of* $\{\mathbf{x}, y\} \in \{\mathbf{X}, Y\}$.

Under this condition the individuals maximization problem has only one stationary point, and the utility goes down while moving away from this optimum. The equations (2.7)- (2.9) can then only describe the unique optimum for all agents. Although this result sounds quite trivial, it has a very interesting application. Because equation (2.9) is a function of the tax schedule, the planner can set the taxes in such a way that this condition is met. The implementation based on the principle of taxation shows a very extreme version of a tax schedule that satisfies this constraint, but quite generally less restrictive tax schedules will suffice.

Note that both of these remarks discuss sufficiency results. To the best of our knowledge no necessary and sufficient conditions for finding the global maximum are known. however, most applications satisfy the conditions in the first remark, or can be made to satisfy the second. In order not to repeat the discussion, we will assume that the maximum identified is the global maximum in the rest of the paper.

The Mirrleesian Tax System

As a corollary to proposition 2.4 we show when the Mirrleesian tax system can implement a second-best allocation.

Corollary 2.8. *Under uni-dimensional heterogeneity,* $p = 1$*, when an allocation i.) lies on the Pareto frontier, or ii.) is optimal to a welfarist planner, and the Mirrleesian tax* *system, as defined in section 3.4, does not contain an internal maximum in any vector subspace of* $\mathbf{X}^*$*, the Mirrleesian tax system can implement the second-best allocation.*

Proof. In the Mirrleesian tax system tax rates are equated to wedges such that all conditions in i.) proposition 2.4, respectively ii.) corollary 2.5 are satisfied. □

Corollary 2.8 is a simple extension of proposition 2.4 and corollary 2.5. The Mirrleesian tax system, which is only defined under uni-dimensional heterogeneity, equates taxes to wedges. Hence, it can implement any incentive-compatible allocation on the Pareto frontier, as well as the second-best allocation of a welfarist planner.

This corollary is important primarily because the Mirrleesian tax system is separable. Hence, our result show that a separable tax system can implement the second-best allocation of a welfarist planner under uni-dimensional heterogeneity of agents. In static models this implies, for example, that the tax rate on capital income is independent of labor income and vice versa. In non-stochastic dynamic models, it implies that an implementation exists with tax rates that are independent over time, such that the government does not have to keep records over time to attain its optimal allocation.

2.5.3 Surjective allocations

There is an important second class of maximization problems in which equating taxes to wedges will always implement the second-best allocation. If the optimal allocation is surjective, such that all possible bundles of goods that are available in the market are assigned to at least one type in the second-best allocation, then only one tax system can satisfy first-order implementability conditions (2.7) and (2.8), and this tax system implements the second-best allocation. Proposition 2.9 formalizes this result.

Proposition 2.9. *If the mapping* $\mathbf{x}^*(\mathbf{n})$ *is surjective, then the tax system is fully described by equations (2.7) and (2.8), and this is the unique differentiable tax schedule that implements the second-best allocation.*

Proof. The proof can be found in the appendix. □

Note that surjectiveness of the mapping $\mathbf{x}^*(\mathbf{n})$ is a relatively strict requirement. It requires that every possible bundle of $\mathbf{x} \in \mathbf{X}$ is assigned to at least one type. In this situation every choice in the market corresponds to the choice of a type in the direct mechanism. Since all types prefer their own bundle over the bundles assigned to other types and all bundles are assigned to a type, it follows that all types prefer their bundle above any other bundle in the economy.

Note that the tax system described in proposition 2.9 is unique. Wedges are defined for every possible value of $\mathbf{x} \in \mathbf{X}$, such that only one implementing tax system exists. Moreover, since surjectiveness implies that $\mathbf{X}^* = \mathbf{X}$ the principle of taxation does not restrict the choice set of agents in this particular case.

The allocation derived in Mirrlees (1971) is an example of a surjective allocation, provided ability is continuously distributed in $\mathcal{R}_+$, and the second-best allocation assigns all positive gross income levels to at least one type.[13] In this case, the function $x^*(n)$, mapping ability to gross income, is surjective.

2.5.4 Implications for implementation in the existing literature

Our paper provides new insights about tax systems for a number of articles in the literature. In most models considered in the existing literature the planner is welfarist, there are no externalities, agents are heterogeneous in only one dimension and the optimal resulting tax system is monotonic or convex (see e.g. Mirrlees, 1971, Mirrlees, 1976, Atkinson and Stiglitz, 1976, Saez, 2001,Bovenberg and Jacobs, 2005, Golosov et al., 2013). Additionally, a number of articles study the implementation of allocations on the Pareto frontier, without specifically assuming a benevolent planner (e.g. Werning, 2007 and Brendon, 2013). By proposition 2.4 and its corollary 2.8 the second-best allocation in these papers can be implemented by a Mirrleesian tax system. In all aforementioned articles, except Golosov et al. (2013), this tax system is also suggested. Golosov et al. (2013) proposes to implement the allocation through a non-separable tax system that also equates taxes to wedges. Proposition 2.4 proves that this tax system implements the allocation derived in the article as well.

A growing literature considers optimal taxation under multi-dimensional heterogeneity of agents, where the planner is welfarist, there are no externalities, and the tax system that naturally results from equating taxes to wedges is monotonic or convex (see e.g. Mirrlees, 1976, Choné and Laroque, 2010, Renes and Zoutman, 2014a, Rothschild and Scheuer, 2014 and Jacquet and Lehmann, 2014). Proposition 2.4 also applies to these models, such that any tax system that equates taxes to wedges implements the allocation derived in these articles.[14]

A number of articles study optimal non-linear taxation in the presence of a non-welfarist planner (see e.g. Fleurbaey and Maniquet, 2006, Fleurbaey and Maniquet, 2007, Jacquet and Van de Gaer, 2011 and Gerritsen, 2013) under uni-dimensional heterogeneity

[13]We implicitly assume agents can also choose their gross income level in $\mathcal{R}_+$.

[14]Note that Rothschild and Scheuer (2014) also extend their analysis by introducing an externality. For this extension proposition 2.4 does not apply.

in earning ability. In each of these articles the agents only choose how much income to earn and how much to consume. Therefore, provided the distribution of earnings ability is continuous and unbounded, the allocation is likely to be surjective. The surjectiveness of the allocation can easily be verified by checking whether all non-negative income levels are assigned, as is usually the case. In that case proposition 2.9 shows that there is a unique tax system that implements the second-best allocation by equating the tax to the wedge.

2.6 Concluding Remarks

The results presented in this paper cover the relatively understudied second step of tax design for a general class of models. Propositions 2.4 and 2.9 show that a relatively simple class of tax systems, that can be described by the optimal wedges in the second-best allocation, can implement the optimal allocation found in most problems studied in the existing literature. In these models implementation can proceed through any tax system that equates taxes to wedges. In addition, lemma 2.3 provides the conditions that need to be checked if the problem does not fit one of the two classes outlined in our propositions. Importantly, our results are not restricted to the often-studied case of uni-dimensional heterogeneity of agents, but are also directly applicable to the multi-dimensional problem studied in the companion paper Renes and Zoutman (2014a).

Proposition 2.4 highlights a unique feature of the Mirrleesian optimal tax model. Unlike the design problem of auctioneers and monopolists, the maximization problem of a welfarist central planner is aligned with that of the agents he faces. In monopoly pricing and auction theory on the other hand, the objectives of the principal and the agents are usually opposed, in the sense that the payoff of the planner increases in the payment made by the agents. An increase in a monopolist's profits (at fixed quantities) automatically comes at the expense of the consumers. As such, implementation will generally be more difficult than in the Mirrleesian taxation literature. As a result, our results indicate that the planner in these settings generally needs to restrict the choice space of its agents by for example bundling goods or disallowing certain bids (see also Armstrong, 1996, Renes, 2011, Rochet and Choné, 1998).

The alignment between agents and planner allows a relatively broad class of tax systems to implement the second best of a welfarist planner. In many cases the planner can let the agents maximize their utility subject to a budget constraint, without further restriction in the choice space, irrespective of the actual utility function of individuals.

This result is important because restricting the choice space of agents in a real-world economy is generally very complex, and in most cases unrealistic and undesirable in a market economy. However, our results may also help explain cases where the price mechanism fails and the government has to restrict the choice space of agents. This may for instance be the case in settings where externalities play an important role.

Future work could focus on the possibility to extend this work to dynamically stochastic settings. Additionally, finding implementing tax systems that allow individual freedom of choice in cases when equating taxes to wedges does not lead to the optimal allocation is of first order importance to the normative public finance literature that deals with issues like externalities and non-welfarist motives.

Chapter 3

As Easy as ABC? Multi-dimensional Screening in Public Finance[1]

3.1 Introduction

How should a government combine taxes on labor income with healthcare subsidies? What is the relation between capital and labor income taxes? When should housing subsidies depend on wealth and income?

The optimal interplay between tax instruments crucially depends on the relevant dimensions in which agents differ, such as earnings ability, wealth, and health. However, up to now, the literature on optimal taxation has almost exclusively focused on models where agents differ in only one dimension, namely earnings ability (e.g. Mirrlees (1971), Diamond (1998), Saez (2001), Bovenberg and Jacobs (2005), Golosov et al. (2013)), and on models where agents differ in various dimensions but the government can only tax labor income non-linearly (e.g. Saez (2002a), Saez (2002b), Choné and Laroque (2010), Jacquet et al. (2013), Rothschild and Scheuer (2014), Jacquet and Lehmann (2014)). Therefore, the current literature can only guide policy makers on the optimal relationship between these non-linear tax and subsidy instruments under the extreme assumption that the difference between agents can be expressed in a single parameter.

[1]This chapter is based on Renes and Zoutman (2014a). We would like to thank Felix Bierbrauer, Eva Gavrilova, Aart Gerritsen, Yasushi Iwamoto, Bas Jacobs, Laurence Jacquet, Etienne Lehmann, John Morgan, Dominik Sachs, Dirk Schindler, Dana Sisak, Bauke Visser, Casper de Vries and Hendrik Vrijburg for useful suggestions and comments on an earlier version of this paper. Furthermore, this paper benefited from comments and suggestions made by participants at the 2011 Nake Conference, Utrecht, the 2013 CESifo Area Conference on Public Economics, Munich, the 69th IIPF Conference, Taormina; ; seminar participants at the Erasmus School of Economics, the Norwegian University of Science and Technology, the Norwegian School of Economics, the University of Konstanz and the Centre for European Economic Research. All remaining errors are our own.

The current literature on optimal taxation is based on the seminal work of Mirrlees (1971, 1976). These papers characterize the welfare maximizing allocation in a setting with individuals that differ in earnings ability alone. In two very influential contributions Diamond (1998) and Saez (2001) rewrite the solution of the model into a much more intuitive form known as the ABC-formula for optimal taxation. This formula describes the optimal wedge between the marginal rates of transformation and the marginal rates of substitution (further: wedge) as a function of behavioral elasticities and the distribution of income. This ABC-formula facilitates an intuitive explanation of the optimal wedges in the second-best, and serves as a convenient way to approach data. Together these papers have provided the basis for a very fruitful line of research in optimal non-linear redistributive policies.

However successful this approach has been, economists have long recognized that it is necessary to make more realistic assumptions and allow for multi-dimensional heterogeneity in models of optimal taxation (see a.o. Sandmo (1993); Saez (2002a); Judd and Su (2006); Lockwood and Weinzierl (2012)). Furthermore, the urgency to extend our models has increased with recent policy discussions on health care subsidies in the US, that followed the introduction of the affordable care act, as well as discussions on the taxation of capital, inspired by the best-selling book Piketty (2014). In this chapter we extend the Mirrleesian model of taxation by letting agents differ in a vector of characteristics and make a vector of choices, to allow for a richer theoretical discussion of the problem of redistribution.

In our model agents differ in $p \geq 1$ hidden characteristics, such as ability, health-status, initial wealth and/or patience. The agents make k observable continuous choices pertaining for instance, labor income, consumption of healthcare products and savings. Additionally, they choose how much to consume of an untaxed numeraire good. We will often refer to the choice variables as goods, although they can be either inputs or outputs in the production process. We assume each of the k goods can be taxed non-linearly.[2] In order to facilitate full revelation we make two assumptions. First, we assume preferences allow the revelation of all characteristics given the proper mechanism.[3] Second, we assume $k \geq p$, the number of choice variables is larger than the number of hidden characteristics.

[2]It may be possible to extend our model to allow for choice variables that can only be taxed at a linear rate, as was done in the case of uni-dimensional heterogeneity of agents in e.g. Mirrlees, 1976, Christiansen, 1984, Boadway and Jacobs, forthcoming. However, this would severely complicate our analysis, as well as the interpretation of the optimal tax formulas, and is hence beyond the scope of this paper.

[3]In particular, we exclude characteristics that do not influence the preference of any choice variable, and multiple characteristics that influence the preference of only one choice variable. Such characteristics are fundamentally non-revealable in any mechanism.

These assumptions allow types to be revealed in the market mechanism, which allows us to relate optimal policy to observable choice variables in the market. We treat the problem faced by the central planner as a multi-dimensional screening problem. If all relevant characteristics of all agents were known to the central planner, the Second Welfare Theorem would imply the planner could select any efficient allocation and implement it through a schedule of individualized lump-sum taxes. However, the planner cannot observe the type of each individual directly, so he has to incentivize each individual to reveal his hidden and multidimensional type. The distortions created by the tax system are the tools used by the planner to gain information about the type of each individual. This information can then be used by the planner to redistribute from one type to another. By reinterpreting the optimal tax problem as a screening or information problem in this way, we can use insights from the multi-dimensional screening literature (most notably McAfee and McMillan, 1988, Armstrong, 1996, and Rochet and Choné, 1998) and apply them to optimal taxation. This reinterpretation is not completely innocuous. The screening literature focuses on direct revelation mechanisms and not on the optimal tax system we want to study. To be able to discuss optimal tax systems, we have to determine how to find a decentralizing mechanism, a collection of tax rates in this case, that implements a given second-best allocation in the market. In chapter 2 we show that the design of the implementing (tax-)mechanism can be very complicated. However, as proposition 2.4 shows, if the government maximizes a welfarist social welfare function and there are no externalities, the government can implement the second-best allocation by a tax system that equates the marginal tax rates to the optimal wedges in the second-best allocation. In this paper we (therefore) assume that the government is welfarist and there are no externalities. Proposition 2.4 guarantees that the optimal wedges found through the direct mechanism in this paper describe all relevant aspects of the optimal tax schedule in our setting. This allows us to side-step the implementation question in this paper, while still describing the characteristics of the optimal tax-schedule.

We compare our characterization of the second-best allocation to the second-best in a uni-dimensional setting by establishing whether well-known results in the uni-dimensional setting hold in our more general model. Our first proposition shows that the optimal wedge can be characterized by a generalized version of Diamond's (1998) and Saez' (2001) *ABC*-formula with an additive structure over the characteristics. The optimal wedge on each good increases in A, the quality of the signal obtained from observing the good: how much does a specific choice reveal about the hidden type. In the standard case with uni-dimensional heterogeneity in earnings ability, an agent's gross labor earnings reveals more about his earnings ability if his elasticity of labor supply is smaller. Hence, in our interpretation, the optimal tax rate on labor income is decreasing in the labor supply elasticity, because as the labor supply elasticity increases, the information gained from distorting labor earnings decreases. In addition, we find that the optimal wedge increases in B, the redistributive benefit of marginally distorting the price of the good, and decreases in C, the size of the tax base for which the marginal choice is distorted. The two latter properties have already been established in the uni-dimensional model. They scale the optimal distortion for welfare and efficiency reasons respectively.

A corollary to proposition 2.4 shows that an optimal wedge on a good is zero if this good does not reveal information about any of the hidden characteristics. This corollary implies the Atkinson-Stiglitz (A-S) theorem in case of uni-dimensional heterogeneity and weak separability of the utility function. In our interpretation, if disutility of labor is the only aspect of utility that is not separable from type, the labor choice is the best signal of the underlying type. Indirect taxation yields no extra information and is thus not optimal.

The corollary also immediately implies the A-S theorem does not generalize to multi-dimensional heterogeneity. With at least two types of heterogeneity, a single good can never extract all available information. If the planner wants to separate agents in the health and ability dimensions simultaneously, he will have to observe and distort at least two choices. For instance by taxing both labor income, as well as a good that reveals the health status of the agent, such as consumption of health care products.

Mirrlees (1976) shows that if agents are heterogeneous in only one dimension, the optimal wedge on each good can be written as a function of only that good. In case of multi-dimensional heterogeneity such separable wedges are practically impossible. We show that in general, in order to facilitate full revelation of the p underlying characteristics, the marginal tax on each choice is a function of p choice variables.[4] Such interdependencies between distortions on separate choices are very common in the stochastic dynamic models of the New Dynamic Public Finance (NDPF).[5] These models extend the Mirrleesian framework to a setting where the type of the agents evolves stochastically over time. In the NDPF the tax on labor income in period t may depend on the labor income earned in periods prior to t (see e.g. Kocherlakota, 2005). Since in these models the history of play contains information about the preferences of the agents, it forms a natural extension to the type space. We show that the interdependencies found in the NDPF models can be replicated in a deterministic Mirrleesian public finance model, provided the type space is

[4]This result describes the general case. Special cases may exist where the wedge can be written in a simpler form.

[5]See Golosov et al. (2007) for an overview

multi-dimensional. This suggests that the intertemporal interdependencies in the optimal tax-schedule in NDPF models may stem from the multi-dimensionality of the type space, rather than from the stochastics.

We also derive a generalization of the no-distortion at the top result (see e.g. Sadka, 1976 and Seade, 1977). As in the uni-dimensional case, the optimal wedge at the extreme points of of the type-distribution are zero. If a type exists that has extreme values for all characteristics, his optimal marginal wedge on all choices equals zero.[6] Intuitively, since there are no "more extreme" types, setting a wedge to separate out more extreme types yields no information to the planner. Hence, for any marginal distortion the efficiency cost of the distortion is higher than the welfare gain at the extreme points of the distribution. Note that unlike in the uni-dimensional case, in the multi-dimensional case these types do not necessarily exist. For instance, the healthiest person in the economy may not be the richest person in the economy. In that case, the healthiest person may face a positive wedge on his labor income, whereas the richest person may face a positive wedge on his consumption of healthcare products.

We overcome the technical complexities of deriving the second-best allocation under multi-dimensional heterogeneity by using a first-order approach. That is, we derive the optimal allocation in a relaxed problem that takes the first-order incentive constraints into account, while assuming the second-order incentive constraints are met in optimum.[7] This approach has become the standard in the optimal taxation literature with uni-dimensional heterogeneity. It is well-known that solutions to multi-dimensional screening problems obtained by the first-order approach consistently violate second-order incentive constraints at the bottom of the type space if participation constraints are binding (see e.g. Armstrong, 1996, Rochet and Choné, 1998). Intuitively, if a principal tries to extract all of the rents to private information from he bottom types, they will simply stop participating. Therefore in the second-best allocation these types are bunched together on the outside option. However, models of optimal taxation typically do not feature binding participation constraints because it is assumed to be too costly to leave the jurisdiction. Hence, there is no inherent conflict between participation constraints and incentive constraints. The normal arguments that guarantee bunching in the second-best allocation, therefore, do not apply. Although we cannot formally prove that bunching never occurs in our model, we show in section 3.8 that if separation of types and bunching occur simultaneously, separation will occur in a single convex subset extending from "the top" of the type

[6]This result is reminiscent of the theorem derived in Golosov et al. (2011) where the wedge at the bottom and top is zero if the stochastic process allows agents to be located at the extremes.

[7]An introduction to this technique can be found in Wilson (1996).

space. In this separating set, our solution obtained through the first-order approach still describes the second-best. Hence, even if the optimal allocation exhibits bunching of types at the bottom of the type space, our solution remains valid in the upper-interior part of the type space where full separation of types is optimal.

The rest of this paper is organized as follows. The next section discusses related literature. Section 3.3 introduces the model. Section 3.4 derives the optimal allocation using the first-order approach. Section 3.5 and 3.6 discuss the ABC-formula and our generalization of A-S. Section 3.7 compares our results to results obtained in the NDPF. Section 3.8 discusses the potential bunching problem and thus the validity of the first-order approach and the final section concludes.

3.2 Related Literature

In this paper we rely on the first-order approach to elicit the properties of the second-best allocation. Another approach to keep the model tractable would be to discretize the type space (see e.g. Armstrong and Rochet, 1999 for a user's guide). In a model with discretely distributed types it is possible to (numerically) verify which incentive constraints are binding, such that the optimal allocation can be derived without relying on the first-order conditions allone. The downside of discretizing the distribution is that the optimal wedge can only be determined on a discrete number of points. Moreover, as the number of discretized types increases, the problem becomes less tractable. Because in our model types are continuously distributed it is possible to calculate the wedge for all levels of the choice variables, thereby deriving the entire shape of the optimal tax system.Cremer et al. (2001) show in a setting with discretely distributed earnings ability and wealth endowments that the A-S theorem fails to hold in this setting since the government optimally taxes savings.

A similar result is derived by Saez (2002a), but now in a n a setting with continuously distributed types. In a model where agents are heterogeneous in both earnings ability and preferences the A-S theorem fails when preferences for a particular commodity are positively correlated with earnings ability, or the preference for leisure. In this case the government should optimally tax these commodities at a higher rate. Our corollary 3.5 will show that this result holds more generally. Unfortunately, two strong assumptions make it difficult to use his approach to calculate the entire tax system. First, he assumes that welfare weights are correlated to ability, but uncorrelated to the other hidden characteristics. However, governments are also likely to give higher welfare weights to agents with lower health status and lower wealth endowments. Second, in his model all goods except labor income are taxed linearly. However, modern governments have access to a wider range of non-linear instruments such as the tax on capital income, health care subsidies and education subsidies. Our approach poses no such restrictions and can be used to calculate all optimal non-linear wedges.

Kleven et al. (2009) study the taxation of couples in a setting where partners have different earnings ability. To maintain analytic tractability they assume the primary earner chooses labor supply on the intensive margin while the secondary earner partner chooses on the extensive margin. In our model agents only make intensive-margin choices. We argue that many economic decisions such as savings and consumption choices are more accurately portrayed as choices on the intensive margin. The best solution would be to combine both approaches. for instance by extending our model with extensive-margin decisions, as was done with uni-dimensional heterogeneity in Jacquet et al. (2013). However, we leave this for future research.

Several papers study multi-dimensional screening in a setting where the number of tools available to the planner is smaller than the number of characteristics, $k < p$. In such a setting full separation in the decentralized mechanism is clearly not possible. Pass (2012) shows that quite generally a less direct version of the revelation principle may be applied. If the planner finds out how to bunch individuals of different types, he can integrate out dimensions of the type space until the adjusted dimension of the type space matches the dimension of the choice space. In the reformulated problem the dimensions match and one can treat it as a "normal" mechanism design problem. This method has been successfully applied in Choné and Laroque (2010) in an optimal-tax model where agents choose labor supply and are heterogeneous in both opportunity cost of work and ability, while the planner only uses an income tax. They show the income tax rate may be negative at the bottom of the income distribution if heterogeneity in the opportunity cost of work is relatively important. Jacquet and Lehmann (2014) develop a similar method and show that, because of the integration over the type space, the optimal tax rates identified by an ABC-formula now depend on average behavioral elasticities at each income level. To limit the complexity of our problem we restrict our attention to the case where $k \geq p$ such that we can get full separation and apply the revelation principle directly. However, it may be possible to extend our results to the case where $k < p$ by applying the method developed in Pass (2012) and Jacquet and Lehmann (2014).

Rothschild and Scheuer (2011, 2013, 2014) study the situation where the extra dimensions in heterogeneity relate to differences in the productivity of each individual in different productive sectors. General equilibrium effects exists in this model because individuals can shift effort from one sector to the other. This forces the planner to adjust his optimal tax rates, compared to the standard Mirrleesian tax rates, to reduce inefficient shifting. In our model we do not explicitly model different sectors, but it does nest a model where each individual decides on the intensive margin of effort in different sectors.

In our model we maximize a standard Bergon-Samuelson welfare function, such that social welfare is a weakly concave sum of individual utility levels. In the field of social choice there is a large discussion about the validity of using this welfarist objective in case of preference heterogeneity. A welfarist planner will generally assign different welfare weights to agents with similar ability levels, but different tastes. It is often argued that such differences in welfare weights violate general notions of fairness. As a result, a growing literature is studying the optimal tax schedule on labor income under a variety of non-welfarist objective functions (see e.g. Fleurbaey and Maniquet, 2006; Fleurbaey, 2006; Kanbur et al., 2006; Jacquet and Van de Gaer, 2011; Ooghe and Peichl, 2011). Their results show that optimal policy depends strongly on the choice of the planner's objective function. However, Kaplow and Shavell (2001) show that any allocation that does not maximize the objective function of a welfarist planner, violates the Pareto principle. In addition, Renes and Zoutman (2014b) show that implementation of the second-best allocation can become very complex, unless it is assumed that the planner maximizes a welfarist objective function. Finally, it is far from clear that governments should give equal weight to two agents with similar ability if one is in significantly better health, or has significantly more inherited wealth than the other. Therefore, in our setting, we apply the welfarist approach, it offers a tractable objective under which one can define optimal policy, as well as facilitates the comparison with the uni-dimensional models.

3.3 The Model

In this section we introduce the formal model that will consequently be solved and discussed in the later sections. First, we define the preferences of the agents in the economy and the conditions for incentive compatibility. We then use these conditions as restrictions in the planner's maximization problem. The model closely follows that of the chapter 2 and of section 4 of Mirrlees (1976).

3.3.1 Preferences

The economy is populated by a unit mass of individuals that are characterized by a twice-differentiable utility-function:

$$u(\mathbf{x},y,\mathbf{n}),$$

where $\mathbf{x} \in \mathbf{X} \subseteq \mathcal{R}^k$ denotes a vector of choice variables such as effective labor supply, consumption of health care products and savings.[8] $y \in Y \subseteq \mathcal{R}$ Is an untaxed numeraire commodity. In principle the choice of the numeraire variable has no effect on the optimal allocation (an undifferentiated tax on $\mathbf{x}$ can achieve the same effect as a tax on y). Decision variables $\mathbf{x}$ and y are observable at the individual level, and the social planner can tax all choices in $\mathbf{x}$ non-linearly, but cannot tax y. Since y will act as an untaxed numeraire, whether a good is taxed or subsidized relative to laissez faire can be evaluated by direct comparison to the untaxed y. We assume y is a normal good, such that $u_y > 0, u_{yy} \leq 0$ for any value of $(\mathbf{x},y,\mathbf{n})$, this directly implies non-satiation of the utility function everywhere. throughout the paper we will sometimes refer to the choice variables in $\{\mathbf{x}, y\}$ as goods, even though they can be both inputs and outputs to the production process.

$\mathbf{n} \in \mathbf{N} \subseteq \mathcal{R}^k$ Denotes the type of an individual. Their type is private information to each individual and unobservable to the government. Each element n_j in the type vector $\mathbf{n}$ is referred to as a characteristic. Characteristics may include for instance earnings ability, health status and preference parameters. For technical convenience we assume that the type space, $\mathbf{N}$, is convex. The distribution of $\mathbf{n}$ is given by the twice-differentiable cumulative density function $F(\mathbf{n})$, with $F : \mathbf{N} \to [0,1]$ with probability density function $f(\mathbf{n})$. Both are defined over the closure of $\mathbf{N}$. For technical convenience we assume $f > 0$ in the interior of N.

Note that we do not restrict ourselves to static models: different choices can occur in different periods. However, we do assume that both their type and the direct mechanism are revealed to the individuals before they solve their maximization problem.[9]

To ensure full separation can occur in a market mechanism we need two additional assumptions. First, we assume that $k \geq p \geq 1$, such that there are at least as many decision variables in $\mathbf{x}$ as characteristics in $\mathbf{n}$. Therefore, the choice space is large enough

[8]Note also that the conventional utility representation(See e.g. Mirrlees, 1971, Saez, 2001.) $\tilde{u}(y,l)$ with l denoting labor supply is a special case of our utility representation. If one takes the standard assumption that gross income equals $x_1 = n_1 l$ where n_1 is earnings ability, it can be seen that this utility function can be rewritten into our form: $\tilde{u}(y,l) = \tilde{u}\left(y, \frac{x_1}{n_1}\right) = u(x_1, y, n_1)$

[9]The model with uni-dimensional heterogeneity has been used often to describe a dynamic economy. See Golosov et al. (2013) for a recent example.

to contain all information in the type space. Second, let:

$$\mathbf{s}(\mathbf{x}, y, \mathbf{n}) \equiv -\frac{u_{\mathbf{x}}(\mathbf{x}, y, \mathbf{n})}{u_y(\mathbf{x}, y, \mathbf{n})},$$

denote the vector of shadow prices, such that each element, $\mathbf{s}_i$, denotes the marginal rate of substitution for decision variable x_i with respect to the numeraire y. We assume each characteristic denotes some independent aspect of the individuals, such that no characteristic can be found as a deterministic function of the other characteristics. As a sufficient condition for this independence we assume the Jacobian $\mathbf{s}_{\mathbf{n}}$ is of full rank, p, for any combination $\{\mathbf{x}, y, \mathbf{n}\}$. This assumption excludes the possibility of having characteristics that do not influence marginal preferences and the possibility of having two characteristics that jointly affect the preference of only one choice. An example of the later is a model where individuals differ in their degree of earnings ability and in their opportunity cost of work. The utility cost of providing a unit of effective labor supply is decreasing in ability and increasing in the opportunity cost of work. If both characteristics act only on effective labor supply, it is fundamentally impossible to separate them both in the choice of labor supply. By assuming $\mathbf{s}_{\mathbf{n}}$ is of full rank, we guarantee that there is always a second observable choice which can be used to disentangle the effect of ability and the opportunity cost of work. If, for instance, the planner could also observe the time spend on video games, while the preference for video games increases in the opportunity cost of work, the problem can be solved. In that case the planner can deduce both characteristics by jointly observing labor earnings and the time spend on video games.[10]

For bookkeeping, the Jacobian of first-order derivatives $\phi'(\cdot)$ of any function $\phi(\cdot) : \mathcal{R}^a \to \mathcal{R}^b$, is of dimension $b \times a$, while the second-order derivatives $\phi''(\cdot)$ are of dimension $ab \times a$. For any multi-vector functions $\psi(\mathbf{z}^1, \mathbf{z}^2, \ldots) : \mathcal{R}^{a^1} \times \mathcal{R}^{a^2} \ldots \to \mathcal{R}$ the vector of first-order derivatives ψ_{z^i} are of dimension $a^i \times 1$ and the matrix of second-order derivatives $\psi_{z^i z^j}$ are of dimension $a^i \times a^j$ where the dimension of the matrix follows the order of the subscripts. Superscript T denotes the transpose operator. Vectors and multi-dimensional constructs are denoted in bold, scalars are in normal font.

[10]Choné and Laroque (2010), Pass (2012) and Jacquet and Lehmann (2014) develop a screening model with non-revealable characteristics. They solve this problem by aggregating the characteristics into a set of revealable virtual characteristics. It may be possible to expand our approach using their methodology, but this is beyond the scope of this paper.

3.3.2 Incentive Compatibility

Before we go to the problem faced by the social planner, we need to consider the problem of the individuals in our economy. In particular, we derive conditions under which an allocation is incentive compatible. These incentive compatibility constraints will subsequently be used to solve for the optimal allocation. In a direct mechanism the social planner offers bundles $\{\mathbf{x}^*(\mathbf{m}), y^*(\mathbf{m})\}$ for all $\mathbf{m} \in \mathbf{N}$. Each individual selects a bundle $\{\mathbf{x}(\mathbf{m}), y(\mathbf{m})\}$ by sending a message $\mathbf{m} \in \mathbf{N}$ to the social planner. Function $\mathbf{x}^*$ maps from the message space to the choice-variable space, $\mathbf{x}^* : \mathbf{N} \to \mathbf{X}$ and y^* maps from the message space to the numeraire commodity space, $y^* : \mathbf{N} \to Y$. An allocation $\{\mathbf{x}^*(\mathbf{m}), y^*(\mathbf{m})\}$ is incentive compatible if each individual truthfully reveals all his unobserved characteristics and receives the bundle designed for him.

$$\mathbf{n} = \arg\max_{\mathbf{m}} u\left(\mathbf{x}^*(\mathbf{m}), y^*(\mathbf{m}), \mathbf{n}\right) \quad \forall \quad \mathbf{n} \in \mathbf{N} \tag{3.1}$$

Let:

$$V(\mathbf{n}) \equiv \max_{\mathbf{m}} u\left(\mathbf{x}^*(\mathbf{m}), y^*(\mathbf{m}), \mathbf{n}\right) \tag{3.2}$$

denote the indirect utility function as a function of the characteristics. In an incentive compatible allocation $V(\cdot)$ satisfies:

$$V(\mathbf{n}) = u\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right)$$

This equation simply states that maximized utility equals the utility function under optimal choices. Proposition 2.4 below largely follows Mirrlees (1976) and McAfee and McMillan (1988). It derives the first and second-order conditions that are satisfied in a differentiable incentive compatible allocation on an interior maximum.

Proposition 3.1. *An allocation* $\{\mathbf{x} = \mathbf{x}^*(\mathbf{n}), y = y^*(\mathbf{n})\}$ $\forall \mathbf{n} \in \mathbf{N}$ *is incentive compatible if:*

$$y^{*\prime}(\mathbf{n}) = \mathbf{s}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right)^T \mathbf{x}^{*\prime}(\mathbf{n}), \tag{3.3}$$

$$\mathbf{x}^{*\prime}(\mathbf{n})^T \mathbf{s}_{\mathbf{n}} \lessdot 0, \tag{3.4}$$

where $\lessdot 0$*, signifies negative semi-definiteness of the matrix.*

Through the envelope theorem a fully equivalent set of conditions can be derived:

$$V'(\mathbf{n}) = u_{\mathbf{n}}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right)^T, \tag{3.5}$$

$$u_{\mathbf{nn}}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right) - V''(\mathbf{n}) \lessdot 0. \tag{3.6}$$

Proof. The proof can be found in the appendix. □

Equation (3.3) states that an individual should be indifferent between truth telling and mimicking at the margin for all characteristics. For each row j the left-hand side of the equation denotes the gain in y as a consequence of marginally changing the reported characteristic n_j. The right-hand side denotes the utility loss in $\mathbf{x}$ measured in units of y for the same change. Therefore, equation (3.3) states that in equilibrium the marginal cost of mimicking equals the marginal benefits for all characteristics. Equation (3.4) is the usual second-order condition as derived by Mirrlees (1976). If the marginal rate of substitution for decision variable x_i is increasing (decreasing) in characteristic n_j, $(s_i)_{n_j} > 0$ $((s_i)_{n_j} < 0)$, and the allocated amount of the good is also increasing (decreasing) in the characteristic, $(x_i^*)'_{n_j} > 0$ $((x_i^*)'_{n_j} < 0)$, the allocation induces self-selection. It implies higher (lower) quantities of the good are assigned to people with a stronger (weaker) preference for the good.

Equations (3.5) and (3.6) are fully equivalent formulations of the same incentive constraints. They are derived through the envelope theorem. Although their explanation is less intuitive, they are extremely convenient mathematical expressions in the derivations in subsequent sections.

Note that equation (3.4) combined with the assumption that $\mathbf{s}_n$ is of rank p, implies that full separation occurs on the market. This is shown in the next lemma:

Lemma 3.2. *If the allocation satisfies (3.4), $\mathbf{s}_n$ is of full rank and $k \geq p$, then all characteristics are revealed through the bundles chosen by the agents in the direct revelation mechanism.*

Proof. Note that (3.4) can only be satisfied if the product $\mathbf{x}^{*\prime}(\mathbf{n})^T \mathbf{s_n}$ is definite, and hence of full rank, p. Since in a matrix product $rank\,(AB) \leq \min\,(rank\,(A)\,, rank\,(B))$, it follows that (3.4) can only be satisfied if the Jacobian of the allocation, $\mathbf{x}^{*\prime}(\mathbf{n})^T$, is also of full rank p for all values $\mathbf{n} \in \mathbf{N}$. Since $k \geq p$ it follows that the allocation is locally invertible around point $\mathbf{n}$ for all $\mathbf{n} \in \mathbf{N}$. Hence, at least one inverse function from the image of the allocation function to the type space exists: $(x^*)^{\leftarrow} : \mathbf{X}^* \rightarrow \mathbf{N}$, where $\mathbf{X}^*$ denotes the image or range of the allocation function. It follows that by observing the bundle chosen by the agent, one can deduce all his characteristics. □

By lemma 3.2 if the second order incentive constraints (3.4) are satisfied, it follows that the type of the agent can be deduced by observing all his choice variables. This is convenient for our analysis, since it allows us to relate optimal policy to observable choices and underlying characteristics.

3.4 The Second Best Allocation: A First-Order Approach

Now that we have established conditions for incentive-compatibility we can turn our attention to the social planner. We solve the social planner's problem using a direct mechanism. We will use the first-order approach, and assume that the second-order incentive compatibility conditions are met in the optimum. This can be verified ex-post by checking whether equation (3.4) or, equivalently, equation (3.6) is satisfied. We will return to the problem of violations of the second-order constraints in section 3.8.

3.4.1 The government

The social planner is assumed to maximize a concave sum of individuals' utility:

$$SW = \int_{\mathbf{N}} W\left(u\left(\mathbf{x}, y, \mathbf{n}\right)\right) dF\left(\mathbf{n}\right), \tag{3.7}$$

$$W' > 0, W'' \leq 0, \tag{3.8}$$

where $W(\cdot)$ is a Bergson-Samuelson welfare function.[11] We assume the social planner commits to the allocation he offers, he cannot alter the allocation after types are revealed.[12] Redistribution is considered welfare increasing because of (at least) one of two reasons. First, concavity in the utility functions of the individuals would imply that individuals with higher income have a lower marginal utility of income. Second, $W'' < 0$ would imply the social planner gives a higher welfare weight to individuals with lower utility.

The social planner is bound by the economy's resource constraint:

$$\int_{\mathbf{N}} y\left(\mathbf{n}\right) dF\left(\mathbf{n}\right) + R \leq \int_{\mathbf{N}} q\left(\mathbf{x}\left(\mathbf{n}\right)\right) dF\left(\mathbf{n}\right), \tag{3.9}$$

where R denotes exogenous government expenditure and $q(\cdot)$ is the economy's production of y as a function of the decision variables in $\mathbf{x}$. A partial derivative q_{x_i} may be either positive or negative depending on whether choice variable x_i is an input to, or an output of the production process. We assume the production technology exhibits diminishing marginal returns, i.e. $q_{x_i x_i} \leq 0$ for all goods x_i, to guarantee that an interior solution will be reached in laissez faire.

[11] As we show in proposition 2.4 this assumption ensures that the second-best allocation can be implemented by equating taxes to wedges.

[12] See Roberts (1984) for a discussion on the issue of commitment.

3.4.2 First-order conditions

In the first-order approach the social planner maximizes social welfare subject to the first-order incentive compatibility constraint, (3.5), the feasibility constraint, (3.9), and a constraint that allows us to substitute out the utility function for the indirect utility function:

$$\begin{aligned} & \max_{V(\mathbf{n}),\mathbf{x}^*(\mathbf{n}),y^*(\mathbf{n})} \int_{\mathbf{N}} W\left(V\left(\mathbf{n}\right)\right) dF\left(\mathbf{n}\right), \quad s.t. & (3.10) \\ 0 & \geq R + \int_{\mathbf{N}} \left(y^*(\mathbf{n}) - q(\mathbf{x}^*\left(\mathbf{n}\right))\right) dF\left(\mathbf{n}\right), \\ V'\left(\mathbf{n}\right) & = u_{\mathbf{n}}\left(\mathbf{x}^*\left(\mathbf{n}\right), y^*\left(\mathbf{n}\right), \mathbf{n}\right)^T, & (3.11) \\ V\left(\mathbf{n}\right) & = u\left(\mathbf{x}^*\left(\mathbf{n}\right), y^*\left(\mathbf{n}\right), \mathbf{n}\right), & (3.12) \end{aligned}$$

where maximized utility $V(\mathbf{n})$ is explicitly modeled as a choice variable. The Lagrangian to this problem is given by:

$$\mathcal{L}= \int_{\mathbf{N}} \left[\left(W\left(V\right) - \lambda\left(R + y^* - q(\mathbf{x}^*)\right)\right) f + \theta^T \left(V'^T - u_{\mathbf{n}}\right) + \eta\left(u - V\right)\right] d\mathbf{n},$$

where λ is the Lagrangian multiplier associated with the resource constraint, $\theta(\mathbf{n})$ is a p-column vector of Lagrangian multipliers for the set of local incentive compatibility constraints, and $\eta(\mathbf{n})$ is the Lagrangian multiplier that ensures maximized utility equals the utility function for each type. Note that s, f, F, θ, u and their derivatives depend on $\mathbf{n}$, but for clarity of exposition this notation is suppressed. We let $\partial\mathbf{N}$ denote the boundary of $\mathbf{N}$ and $\mathbf{e}$ the outward unit surface normal vector to the boundary of $\mathbf{N}$. through the divergence theorem (or multi-dimensional integration by parts) we can rewrite the Lagrangian as:

$$\begin{aligned} \mathcal{L} & = \int_{\mathbf{N}} \left[\left(W\left(V\right) - \lambda\left(R + y^* - q(\mathbf{x}^*)\right)\right) f - V \sum_{j=1}^{p} \frac{\partial\theta_j}{\partial n_j} - \theta^T u_{\mathbf{n}} + \eta\left(u - V\right)\right] d\mathbf{n} \\ & \quad + \int_{\partial\mathbf{N}} [V\theta^T\mathbf{e}] d\partial\mathbf{N}. & (3.13) \end{aligned}$$

Assuming the functions V and θ are smooth, this function can be maximized point wise on the interior and boundary of the type space.

On the interior of the type space the first-order conditions with respect to variables $\mathbf{x}$, y and V are:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial y} &= 0 : -\lambda f - u_{y\mathbf{n}}\theta + \eta u_y = 0, && (3.14)\\
\frac{\partial \mathcal{L}}{\partial \mathbf{x}} &= \mathbf{0}_k : \lambda q'^T f - u_{\mathbf{x}\mathbf{n}}\theta + \eta u_{\mathbf{x}} = \mathbf{0}_k, && (3.15)\\
\frac{\partial \mathcal{L}}{\partial V} &= 0 : W' f - \sum_{j=1}^{p} \frac{\partial \theta_j}{\partial n_j} - \eta = 0. && (3.16)
\end{aligned}$$

The next proposition uses these first-order conditions to derive an *ABC*-formula for the optimal wedge in the spirit of Diamond (1998) and Saez (2001).

Proposition 3.3. *The optimal wedge on good i for type $\mathbf{n}$ can be described by the following formula:*

$$\begin{aligned}
\frac{q_{x_i}(\mathbf{x}^*(\mathbf{n})) - s_i(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n})}{s_i(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n})} &= \sum_{j=1}^{p} A_{ij}(\mathbf{n}) B_{ij}(\mathbf{n}) C_{ij}(\mathbf{n}) && (3.17)\\
\forall \quad i &= 1, \ldots, k; \mathbf{n} \in \mathbf{N},
\end{aligned}$$

where:

$$\begin{aligned}
A_{ij}(\mathbf{n}) &\equiv \varepsilon_{x_i n_j}(\mathbf{n}) = -\frac{\partial s_i(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n})}{\partial n_j} \frac{n_j}{s_i(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n})},\\
B_{ij}(\mathbf{n}) &= \theta_j(\mathbf{n}) \frac{u_y(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n})}{\lambda}, && (3.18)\\
C_{ij}(\mathbf{n}) &= \frac{1}{n_j f(\mathbf{n})}.
\end{aligned}$$

Proof. The proof can be found in the appendix. □

Note that proposition 3.3 provides the optimal wedges, but gives no direct information about the optimal tax rate. However, as we show in proposition 2.4, if the allocation is optimal to a welfarist planner and there are no externalities, a tax system that equates taxes to wedges can implement the allocation. As such, the wedges derived above contain all relevant information for an implementing tax system. We postpone discussion of the *ABC*-formula to section 3.5 where we compare it to the optimal wedge formula under uni-dimensional heterogeneity. In the next subsections we use proposition 3.3 to derive some technical aspects about the optimal tax rate at the boundaries of the type space and suggest an algorithm to solve the model.

3.4.3 Boundary Conditions: No Distortion at the Corners

The boundary conditions can be found by differentiating equation (3.13) with respect to $V(\cdot)$ at the boundary of the type space ∂N. Only the final term of equation (3.13) depends on the boundary. Hence, we find:

$$\theta_j\left(\underline{n}_j\right) = \theta_j\left(\overline{n}_j\right) = 0, \tag{3.19}$$

where $\underline{n}_j$ $(\overline{n}_j)$ represents the type that has the lowest (highest) value for characteristic j. Define corner types $\mathbf{n}_{\lrcorner}^{\urcorner}$ as agents that have either highest or lowest values for all their characteristics. In a two-dimensional type space type $\mathbf{n} = (\underline{n}_1, \underline{n}_2)$ and $\mathbf{n} = (\overline{n}_1, \overline{n}_2)$ are obviously corner types, but so are the types that combine the lowest value of n_1 with the highest value of n_2 and vice versa: $\mathbf{n} = (\underline{n}_1, \overline{n}_2)$ and $\mathbf{n} = (\overline{n}_1, \underline{n}_2)$. There are at most 2^p corner types although there may be less, or none at all, depending on whether there is a positive density $f(.)$ at each corner, and on whether or not the distribution is bounded. Corollary 3.4 establishes that for each existing corner type the optimal wedge on all goods equals zero.

Corollary 3.4. *The optimal wedge for any type* $\mathbf{n}_{\lrcorner}^{\urcorner}$ *equals zero if the types exist.*

Proof. From the boundary conditions it follows that $\theta_j\left(\underline{n}_j\right) = \theta_j\left(\overline{n}_j\right) = 0$ for all $j = 1, \ldots, p$. The optimal wedge at the corner types can be found by taking the limit of equation (3.17) if $\mathbf{n}$ goes to an $\mathbf{n}_{\lrcorner}^{\urcorner}$:

$$\lim_{\mathbf{n}\to\mathbf{n}_{\lrcorner}^{\urcorner}} \frac{q_{x_i} - s_i}{s_i} = \lim_{\mathbf{n}\to\mathbf{n}_{\lrcorner}^{\urcorner}} \sum_{j=1}^{p} \varepsilon_{x_i n_j} \frac{u_y \theta_j(\mathbf{n})/\lambda}{n_j f(\mathbf{n})}$$

which equals 0 provided $f\left(\mathbf{n}_{\lrcorner}^{\urcorner}\right)$ does not equal zero, that is, provided the type exists in the economy. □

Corollary 3.4 shows that the no-distortion at the top and the bottom results, as derived in Sadka (1976) and Seade (1977), remain valid in a multi-dimensional framework as long as the type distribution is bounded. Technically, the no-distortion at the corner result follows from the transversality conditions of the optimization problem. Intuitively, if there are no individuals of more extreme type, distorting their choices will not yield any extra information. In terms of our motivating examples, if an individual is the healthiest, most able person around, distorting his choices will not lead to extra information or redistributive benefits, but will come at an efficiency loss. As such, the optimal wedges at the corners of the type distribution must equal zero.

Our results are similar to those of Golosov et al. (2011) who derive the optimal tax rate at the boundary in a model where earnings ability follows a stochastic progress. They show that the optimal tax rate for types that persistently have the highest or lowest skill-realization equals zero provided such types exist.

Mirrlees (1971), Diamond (1998) and Saez (2001) show that the optimal marginal tax rate converges to a constant at the top in the uni-dimensional case, provided the upper tail of ability follows a Pareto distribution. We are not able to derive such a result in the multi-dimensional framework since this requires an explicit solution for all θ_j.

3.4.4 Finding the Optimal Allocation

The last step in the problem of the planner is to find the second-best allocation. Although equation (3.17) gives a useful representation of the wedges, there is no closed-form solution for the optimal allocation. Depending on the specification, deriving the optimal allocation may be a computationally complex process and certainly goes beyond the scope of this paper. However, in this subsection we sketch an algorithm that can solve for the optimal allocation. The solution method described here is largely based on Mirrlees (1976).

First, the set of equations (3.11, 3.12) can be used to solve for $\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\}$ as an (implicit) function of $V'(\mathbf{n}), V(\mathbf{n})$ and $\mathbf{n}$:

$$\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\} = \psi(V'(\mathbf{n}), V(\mathbf{n}), \mathbf{n}).$$

Second, by means of this equation and (3.14) and (3.15) we can solve for $\{\theta(\mathbf{n}), \eta(\mathbf{n})\}$ as an explicit function of $\mathbf{x}^*(\mathbf{n})$, $y^*(\mathbf{n})$ and $\mathbf{n}$, and hence as an (implicit) function of $V'(\mathbf{n}), V(\mathbf{n})$ and $\mathbf{n}$:

$$\{\theta(\mathbf{n}), \eta(\mathbf{n})\} = \hat{\phi}(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})) = \phi(V'(\mathbf{n}), V(\mathbf{n}), \mathbf{n}).$$

Finally, if we substitute this result into the last first-order condition 3.16 it forms a second-order partial differential equation. This differential equation can be integrated numerically under the boundary conditions (3.19) and (3.9). The solution provides us with $V'(\mathbf{n})$ and $V(\mathbf{n})$ which can subsequently be used to find the allocation $\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\}$. From here the optimal wedges can be found by substituting the solution $\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\}$ into the ABC-formula (3.17).

3.5 The ABC Formula

In this section we use the ABC-formula (3.17) to compare the optimal wedge under multi-dimensional heterogeneity to the outcome in the uni-dimensional case. In addition, we use the ABC-formula to revisit the Atkinson-Stiglitz theorem under multi-dimensional heterogeneity in the next section.

As in the uni-dimensional case, the left-hand side of equation (3.17) represents the optimal wedge on good x_i for type $\mathbf{n}$. This distortion is broken down into different factors of interest on the right-hand side.

The A-term is a measure of the informational value of good x_i. Intuitively, if the elasticity $\varepsilon_{x_i n_j}$ is large, it means that the preference for choice x_i strongly increases in characteristic n_j. Hence, x_i is a very strong signal of characteristic n_j, and therefore the optimal wedge is large. Our A-term is more general than the ones derived in Diamond (1998), Saez (2001), and Jacquet et al. (2013) because we use a more general utility function. In Diamond (1998) and Saez (2001) the utility-function is of the form: $u\left(y,l\right) = u\left(y,\frac{x}{n}\right)$, where y is consumption, n is productivity and $x = nl$ is effective labor supply (or labor income). Their A-term is inversely related to the compensated labor supply elasticity. In Jacquet et al. (2013) the assumed utility function is $u^1\left(y\right) + u^2\left(x,n\right)$. Their A-term is also inversely related to the compensated labor supply elasticity. The inverse relation between labor elasticities and marginal tax rates can easily be explained in terms of an information problem. If the labor supply elasticity is large, this means that a small change in the net wage rate leads to a large change in labor income. Therefore, distorting labor income marginally leads to a large behavioral response and an imprecise signal of ability. It follows that the optimal tax rate on labor is decreasing in the labor supply elasticity, since the higher the elasticity the less information is gained from marginally changing the labor choice. As this comparison of the A terms shows, our informational elasticity is therefore related to the inverse of the labor elasticity. The intuition behind the term is the same, whatever elasticity one uses.

The B-term represents the redistributive benefits of distorting choice x_i for characteristic n_j. θ_j Is the Lagrangian multiplier of incentive compatibility constraint n_j. Hence, it represents the welfare cost of separating type $\mathbf{n}$ in characteristic n_j. In equilibrium θ_j should equal the marginal welfare benefit of making the allocation marginally less incentive compatible in choice x_i (i.e. increase the distortion on choice x_i). By multiplying θ_j with u_y and dividing through λ the welfare gain for such a redistribution is expressed in units of the numeraire good. An increase in the marginal welfare benefit of distortion, higher θ_j, logically increases the optimal distortion. Since the shape of this term strongly depends on θ, and thus the solution to a set of PDEs, it is impossible to comment on it more generally.

The C-term is related to the size of the tax base for whom the marginal incentives are distorted by the wedge. The denominator represents the size of the tax base with respect to characteristic n_j. The larger this tax base is, the larger the incidence of the distortion and, hence, the larger the efficiency cost associated with the distortion. Efficiency implies that the size of the distortion is inversely related to the incidence of the distortion, as the C-term clearly shows. In uni-dimensional cases, the C-term is often multiplied by $1 - F(n)$ to make it proportional to the (measurable) inverse hazard-rate of the ability distribution. This is corrected for by dividing the A or B-term through the same factor. In a single-dimensional distribution of types, such fractions have an intuitive interpretation as conditional means. Unfortunately, this interpretation is lost when the type-distribution is multi-dimensional. We therefore present the simplest form of the equation, rather than the version with conditional means.

The largest difference between the uni-dimensional and the multi-dimensional ABC-formula is the need to sum over all characteristics to get the optimal wedge for a good x_i in the latter case. That the optimal distortion with a multi-dimensional type space can be described with a formula that resembles the uni-dimensional description so clearly seems quite remarkable. The similarities are a clear testament to the strength of the ABC-formula as a description of the equilibrium in the Mirrleesian tax model. Especially the additive nature of the wedges over the different dimensions of the type space is surprising. The summation indicates that the optimal wedge on good x_i is the sum of the optimal wedges for each characteristic. For example, if earning labor income is easier for agents with higher ability and for agents with a better health state, the planner can calculate the optimal wedge on labor income by adding the optimal wedge on the basis of redistribution in ability to the optimal wedge of redistribution in health. This additive nature of the wedge could be particularly useful for policy evaluation. In our example, if through the uni-dimensional model of optimal taxation one could determine the optimal wedge on the basis of redistribution in ability, this wedge can serve as a lower-bound on the optimal wedge on labor income for a planner that also wants to redistribute from healthy to sick (provided there is a non-negative correlation between health and income).

The summation also appears to indicate that wedges used to identify types in one dimension can be treated separately from wedges used to identify the other dimensions of the type-space. However, beyond the possibility to set bounds, this appearance is deceiving. Intuitively, it seems unlikely that one can separate out individuals on wealth and income independently of each other. Individuals will treat monetary wealth and monetary income as substitutes in their budgets, thus creating interdependence in choices. This relation will lead to interdependencies between wedges in the second best. In the optimum, the marginal tax on income will have to depend on wealth. Similarly, if the planner wants to redistribute more toward the unhealthy individuals with a given ability, than to similarly able individuals with good health, marginal tax rates will have to depend on income and health. These complexities are discussed next.

3.6 The Atkinson-Stiglitz Theorem

The Atkinson-Stiglitz theorem (Atkinson and Stiglitz, 1976), further A-S theorem, states that indirect taxation is superfluous in a setting where agents are heterogeneous in earnings ability if preferences are homogenous and the utility function is weakly separable in labor. The result has subsequently been generalized in Laroque (2005), Kaplow (2006), Gauthier and Laroque (2009) and Hellwig (2010) to more general utility functions with single dimensional heterogeneity. If the conditions of the A-S theorem (i.e. weak separability and uni-dimensional heterogeneity) are satisfied the government can reach the second-best allocation by only taxing labor income. Equivalently the planner could tax all commodities at an uniform rate. The main application of the theorem is, perhaps, that commodities should be taxed at the same rate over time. That is, the optimal capital-tax rate equals zero.

The next corollary uses our *ABC*-formula to investigate when the optimal wedge on a good equals zero. If all wedges but one are equal to zero, the government can reach second-best with a single tax tool and one could say the A-S theorem applies. The next corollary thus shows us when the A-S theorem holds under multi-dimensional heterogeneity by showing when only a single distortion is necessary to reach second best.

Corollary 3.5. *The optimal wedge on good x_i is zero if $\varepsilon_{x_i n_j} = 0 \quad \forall \quad n_j \in \mathbf{N}$, that is, the optimal wedge is zero if the marginal rate of substitution for x_i does not depend on any characteristic n_j for all types.*

Proof. If the marginal rate of substitution, s_i, is independent of all characteristics n_j, then $\varepsilon_{x_i n_j} = 0 \quad \forall \quad n_j \in \mathbf{N}$, such that $A_{ij} \quad \forall n_j$ are zero and the optimal wedge on x_i is zero by equation (3.17). □

Intuitively, corollary 3.5 shows that the marginal wedge on a good equals zero if the preference for this good is not directly influenced by any characteristic. In that case the choice for this good does not provide any first order information and distorting the choice away from laissez faire is not optimal. It follows immediately from corollary 3.5 that the optimal wedge on all goods except income equals zero in the model of A-S. The assumption of weak separability implies that the marginal rates of substitution for all goods except income are independent of type, such that all $\varepsilon_{x_i n_j}$ except the one on income are zero. Because we have assumed $\mathbf{s_n}$ has rank p, $\varepsilon_{x_i n_j} \neq 0$ for at least p choices in our model. Hence, if all characteristics are independent and revealable, the A-S theorem does not hold under multi-dimensional heterogeneity. Intuitively, a government that wants to redistribute in multiple dimensions, cannot do so by distorting the price of only one good. In fact, under our full revelation assumptions of lemma 3.2 a Tinbergen rule applies. A planner that wants to redistribute over p characteristics will need to distort (at least) p choices.

In the literature many violations of the A-S theorem have been recorded. In Erosa and Gervais (2002) preferences are not weakly separable over time since consumption at old age is a stronger complement to leisure than consumption at a younger age. Therefore, the distortion of the labor income tax is reduced by taxing capital income. In Golosov et al. (2013) capital is optimally taxed if households with higher ability also have higher patience. Boadway and Pestieau (2011) show that the theorem fails if households optimally choose a corner solution. Therefore, if some households are cash-constrained and do not buy all commodities, the tax rate on these commodities should be different to identify the cash constrained households. Farhi and Werning (2010) and Kopczuk (2013) show that the bequest motive may generate a negative externality which can be remedied through the taxation of capital. The argument that is closest to our is derived in Cremer et al. (2001) and Saez (2002a) who show that under bi-dimensional heterogeneity commodity taxation is not superfluous. However, the former result is derived in a setting with discrete types and the latter is derived under the assumption that welfare weights are only correlated with earnings ability and commodity taxes are linear.

Our result adds to this literature by showing generally that the A-S theorem cannot hold under multi-dimensional heterogeneity, provided that all (or at least two) underlying characteristics can be revealed. This has large implications for the evaluation of government policy. According to the A-S theorem we can obtain a second-best allocation if the only government intervention is the taxation of labor income. We show that if a government cares about redistribution in multiple dimensions, such as from healthy to sick and from rich to poor, it needs to distort multiple choices in order to attain the second-best allocation. Therefore, government intervention in many markets, like the health care and the rental market, may be optimal. In addition, the optimal wedge on capital income may be non-zero if households differ in their investment skills, for no other reason than this difference in skill.

3.7 Interdependencies in the Tax System in Mirrleesian and New Dynamic Public Finance

The NDPF, pioneered in Golosov et al. (2003), generalizes the Mirrlees model with unidimensional heterogeneity to a dynamic stochastic setting. For each agent, earnings ability follows a known stochastic process that takes a different value in each period. Each agent's earnings ability is revealed to him at the beginning of the period. Agents therefore do not know what their earnings ability will be in future periods, but they can have expectations about it. The planner does not observe ability, but does observe all choices made by the agents in the economy, and can keep records over time.

One of the most intriguing results in the NDPF is the complexity of the optimal tax system. Kocherlakota (2005), for instance, shows that the optimal tax on labor income in period t may depend on the entire history of labor income up to period t.[13] However, so far there is no clear explanation why interdependencies are present in the NDPF in the first place, or alternatively, why they are typically absent in the classical (static) Mirrleesian public finance. In this section we show that interdependencies in the wedges generally do occur in Mirrleesian public finance models, provided agents are heterogeneous in multiple dimensions. Subsequently, we give our model a dynamic interpretation, and provide a link between the result found in Kocherlakota (2005) and our model.

The next corollary uses proposition 3.3 and lemma 3.2 to show that the optimal wedge on each choice depends on p choices.

Corollary 3.6. *The optimal wedge on good i depends on p choices.*

Proof. By lemma 3.2 there exists an inverted mapping of the allocation $(x^*)^{\leftarrow} : \mathbf{X}^* \to \mathbf{N}$ for all $\mathbf{n} \in \mathbf{N}$. This inverse mapping generally depends on p choice variables, since the allocation is of rank p everywhere on the type space. Therefore, we can write the optimal wedges for any type as a function of p choices. Equation (3.17) shows that the wedges solve a set of partial differential equations that are of rank p as a result of lemma 3.2 (see section 3.4.4), hence each individual wedge can be written as a function of p choices. □

[13] Subsequent papers have made some progress on limiting the number of interdependencies in special cases. In particular, Albanesi and Sleet (2006) show that the intertemporal interdependencies disappear when the stochastic process is iid. In that case, they show the optimal wedge in each period depends on current labor income and wealth, such that only an intratemporal interdependency remains between the optimal wedge on capital and labor income.

If all wedges are functions of p choices, then by construction so are all marginal taxes (see Mirrlees, 1976 and Renes and Zoutman, 2014b). Therefore, a separable tax system where the wedge of each good depends only on consumption of that good will not generally exist if $p > 1$. Intuitively, the choice for each good depends on the budget set and the rest of the choices made. Since each individual will optimally trade consumption of one good for consumption of all other goods, the wedge on each good will in general have to condition on the consumption of all other goods. For instance, if the optimal labor income choice depends on both earnings ability and health, the government needs to condition its tax rate on two observable choices that reveal both characteristics (e.g. labor income and consumption of health care products) to be able to generate full separation.

There is a number of mathematical reasons why the optimal wedge for each good is almost always a function of all underlying characteristics. First, most probability density functions, $f(\mathbf{n})$, are a function of all characteristics such that the C term in equation (3.17) depends on all characteristics. In addition, the partial differential equations (3.16), which solve for the θ's, are a function of indirect utility. In equilibrium, indirect utility always has to be a function of all $\mathbf{n}$ in order to fulfill the incentive compatibility constraint (3.11). The solution to the set of partial differential equations in equation (3.17) therefore depends on all equations. It may be possible to construct special cases where the optimal wedge on each good depends on only one characteristic by, for instance, assuming (weak) separability of preferences, a uniform type distribution and quasi-linearity of the utility functions. In general, however, we expect any optimal tax system that redistributes in multiple dimensions to have an amount of interdependencies equal to the dimensionality of the problem.

3.7.1 A Dynamic Interpretation

Consider a T-period economy and assume that agents' abilities evolve over time. Denote ability in period t by n_t and assume the agents know their entire (future) history of ability levels, the vector $\mathbf{n} = [n_1, \ldots, n_T]$ at the beginning of the first period. Clearly, this is a strong assumption on the information available to the agents of the model. In contrast, in NDPF models agents only learn their ability in each period at the beginning of that period, such that their knowledge of future types is limited to knowledge about the stochastic process. Note, however, that this model can be seen as a special case of a NDPF model where ability evolves according to a fully deterministic process known to the agents, but not the planner. The government cannot observe the individual ability levels but is aware of the probability density function of the cross section $f(\mathbf{n})$, and cumulative density function $F(\mathbf{n})$ and both are time independent. In addition, assume for simplicity that each agent makes one independent choice each period, the amount of labor income he earns, x_t. Our model can be used to calculate the welfare-optimizing wedge on each period's labor income x_t. By proposition 3.3 the optimal wedge is given by equation (3.17). From corollary 3.6 it follows that the optimal wedge in each period can be written as a function of the entire history of labor income (including future periods in this case). Intuitively, the optimal wedges depend on terms A, how much information a choice reveals, B, the redistributive value of the wedge, and C, the incidence of marginal distortions. Each of these terms may depend on the entire vector $\mathbf{n}$. In turn, in order to reveal the entire vector the planner needs to observe all choices in $\mathbf{x}$, the entire history of labor income. Due to information restrictions the taxes in NDPF models are generally not conditioned on future income. Since neither the planner nor the individuals know future types, all they can do is calculate expectations. These expectations will depend on the stochastic process (public knowledge) and the realizations of this process, i.e. the history up until that moment. On any truthful path an optimal tax schedule that can non-linearly use the history up until that time, thus envelopes all available knowledge about future types. Adding terms that explicitly condition on unknown future realizations thus add no value.

It may perhaps seem surprising that Kocherlakota's result can be replicated so easily in a deterministic setting. However, as was noted in Pavan et al. (2011), in a deterministic model agents can plan their entire choice vector $\mathbf{x}$ in the first period, with full knowledge about their type $\mathbf{n}$. In a stochastic setting information concerning the type is revealed over time. As such, a stochastic model may allow for less (profitable) deviations than a deterministic model. At any time agents can see what choices they still have to make, such that their choice-vector is known. However, in a stochastic setting agents always have weakly less information than in the deterministic setting. This implies the incentive constraints are more binding in our model than in the NDPF. Hence, if intertemporal interdependencies are optimal in a stochastic setting, they are likely also optimal in a deterministic setting. It also follows that interdependencies in the tax system can stem from the multi-dimensionality in the type space, irrespective of the (dynamic) stochastic processes assumed.

The practical implications of this result are less clear. A literal reading of the model described in this subsection would suggest the government should build intertemporal interdependencies in the tax system. One could let the labor income tax increase in previous income, for instance. However, it is unclear how much welfare is gained by introducing these interdependencies. Simulations on a NDPF model in Farhi and Werning (2013) show that much of the welfare gain of optimal taxation can be obtained by tax systems that do not exhibit any interdependencies. While implementation and administration costs of these interdependencies can be quite high.

3.8 Bunching

So far we have worked from a relaxed problem that assumes away second-order incentive constraints. We have therefore ignored possible violations of the second-order incentive compatibility conditions, (3.4). In most multi-dimensional screening problems the second-best allocation contains bunching at the lower end of the type distribution. Bunching occurs if the solution to the first-order approach violates the second order conditions (3.4) on part of the type space. In this case, certain types would prefer the bundle of another type over the one assigned to them. In our examples wealthy, highly able individuals might prefer the bundle the planner intents for low ability individual, and consume a lot of leisure.

In this section we address this issue by showing what happens to the solution if bunching occurs. The next proposition shows that if bunching occurs, it occurs at the bottom of the type space. Denote by $\mathbf{N}_B$ the bunching partition of the type space, and by $\mathbf{N}_S$ its complement, the separating region, where full separating is optimal. Then by definition $\mathbf{N}_B \cup \mathbf{N}_S = \mathbf{N}$. The next proposition shows that $\mathbf{N}_S$ is a single convex set that extends from the upper boundary of the type space.

Proposition 3.7. *If $\mathbf{N}_S$ exists, it is a single convex set that extends from the upper boundary of the type space.*

Proof. The proof can be found in the appendix. □

Intuitively, the proof of proposition 3.7 follows from the fact that incentive problems are more likely to occur at the bottom of the type space. The planner wants to extract as much informational rents from types at the top of the type space as possible, such that it can redistribute these rents to the bottom types. Therefore, bunching at the top of the type space is particularly costly to the planner. To facilitate full separation of the types at the top, it may be optimal to bunch together types at the bottom of the distribution. In screening problems with binding participation constraints there is always bunching at the bottom of the type space. This was proven in Zheng (2000) for auctions, Armstrong (1996) for non-linear multi-product monopoly pricing and Rochet and Choné (1998) for general screening problems. In all cases, an attempt to extract all rents from agents at the bottom of the type space will inevitably lead to non-participation of some of these agents, and thus a non-participation bunch at the bottom.

Although there is no participation constraint in our model, we cannot exclude the possibility that the first-order approach will violate second-order conditions on part of the type space. However, by proposition 3.7, if separation is optimal in a partition of the type space $\mathbf{N}_S$, this partition is convex. Lemma 3.2 shows that our solution still describes the optimal wedge in the partition of the type space where full separation occurs. It follows, that the optimal wedges in the separating partition are still described by proposition 3.3. Hence, even if the first-order approach does not yield the optimal wedge for every type, this does not necessarily invalidate our approach.

3.9 Concluding Remarks

Although significant progress has been made in multi-dimensional mechanism design, the equilibrium in a multi-dimensional Mirrleesian optimal tax model had so far not been characterized. In this paper we characterize it and show some similarities and differences with the uni-dimensional Mirrleesian model. Furthermore we show how the equilibrium relates to the stochastic dynamic NDPF models and the wider class of multi-dimensional screening models.

Our model can be used to study the relationship between several tax tools. Our characterization of the second-best equilibrium shows that the government should search for consumption patterns that provide as much information on the underlying types as possible. More importantly, the multi-dimensionality in the type space forces the government to make the redistributive taxes depend on several observable choices to separate out different aspects of the hidden types. It might not be optimal to separate out types everywhere in the type space, in which case some bunching occurs at the lower end of the type space. This prescription fits reasonably well with the tax schedules observed in welfare states. The lowest earning individuals get welfare assistance, or income subsidies, creating a bunch at the lower end of the income distribution. Most assistance programs are conditioned on (the absence of) wealth, to make sure that no abuse occurs. This is the kind of interdependencies between underlying characteristics (wealth and ability) our model predicts. Many welfare states also subsidize medical expenses or housing for a large group of people. In theory the government could directly transfer the required money to recipients, rather than paying part of the price. A direct transfer, however, would make it impossible for the government to find out whether or not you are in need of health care, i.e. the government cannot determine your hidden type through a direct transfer, but can do so through the subsidy. We would therefore indeed predict the government uses subsidies that depend on income and expenses (or other observables) rather than direct transfers for differentiated assistance.

The equilibrium in our model depends on solving a set of partial differential equations for which no general solution exists, therefore we can only characterize equilibrium through a set of necessary conditions. These conditions strongly limit the possible outcomes, but can never give a full description of the second-best. The next step in this line of research clearly is to find specific, realistic and relevant settings and simulate the model. This is, however, a difficult step. The multi-dimensional heterogeneity sets strong requirements on the optimization algorithms. In addition, the problem of implementation, which is discussed in chapter 2, might add further difficulties. Implementation can proof especially difficult because implementation on the interior of the separating partition will likely require a different set of instruments than the implementation on the bunching partition. However, once these difficulties have been overcome the model presented in this paper can be used to provide a more precise insight in the optimal relation between the income tax system and the myriad of social schemes like health care subsidies, housing subsidies, and welfare assistance that characterize modern welfare states.

Since this model contains the multi-dimensional type space that is also found in the NDPF without the dynamic stochastics, it could also provide a convenient middle ground between the complex stochastic dynamics in these models, and the known intuitions in the classical Mirrlees model. The problems of joint (or double) deviation and the interdependencies in the optimal wedges that plague NDPF models are, for instance, also prevalent in our setting, but can be traced much more conveniently. These results indicate that at least part of the difficulties in the NDPF literature are due to the structure of hidden information. This shows that we can gain intuition for these tax schedules from multi-dimensional screening models, in particular our model. In fact, the discussion in section 3.7 already suggests that our findings might be generalized to dynamic settings. This would allow an elasticity approach and a new focus on implementation in these models as well.

Chapter 4

Who Dares? Interactive survey evidence on perceptions and risk attitudes in the boardroom.[1]

4.1 Introduction

Individual risk attitudes are an important aspect of economic and finance theory. They are used to explain a.o. the existence of insurance markets, risk premia in asset prices, and option contracts in executive compensation. Risk attitudes have been studied intensively in lab settings using students (see. e.g. Eckel and Grossman (2008)). Most relevant risky decisions, however, are not taken by students but by professionals. Corporate executives in particular are responsible for large investments and risky ventures, almost on a daily basis. Because decision making under risk is such an essential part of the job of an executive and their decisions have a large impact on their companies and employees, studying their decision processes is of first-order importance. Ideally, one would study management decisions related to risk directly and then see the effect on company outcomes. This level of access (and devotion) has proven infeasible so far. The difficulty of bringing together a large group of relevant respondents seems to have caused researchers to look for alternative methods. In behavioral finance, for instance, many authors use demographic attributes or observable behavior of top executives as proxies for their cognitive styles, knowledge base and risk attitudes. These proxies are then used to explain corporate decisions, or stock exchange outcomes to show how the top- management team (TMT) influences the company. While this type of research has uncovered many interesting relationships, it

[1]This chapter is based on De Groot et al. (2012)

leaves many open questions, as it inevitably treats the company as a black box. In this paper we present the results of a new survey methodology that allows researchers to study the decisions of hard to reach groups, like executives, more directly. With this methodology we study how corporate directors view their own and each other's willingness to take risks and compare these beliefs to behavior in realistic investment scenario's.

We use a dynamic website to tailor each survey to the respondent, while guaranteeing anonymity. Each respondent is only asked to respond to vignettes that mimic investment scenarios in areas in which he has expertise and the investment is scaled to the size of decisions he takes in his professional life. This process automates some of the aspects of a structured interview, allowing us to question respondents on relevant aspects of their experience, without ever breaking the guarantee of anonymity, or having to bring large groups of respondents together at any one point in time. This approach allows us to apply some of the methodology of experimental economics, most importantly randomization, pre-structured interaction, exogenously imposed variation in treatments/decisions and anonymity, to a group of relevant decision makers that is near to impossible to get into a lab. With our limited number of respondents we are able to find corroborating evidence for earlier findings from the behavioral finance and upper echelon literature. Furthermore, the exogenous variation allows us to compare between subjects and shed light on the risk attitudes of the different members of the boardroom and their perceptions of each other. We believe that acquiring more quantitative data in this manner provides a viable and valuable method for future research.

The results of our survey show that there is a discrepancy between the perceived willingness to take risk of the average CEO, CFO and Non Executive Director (NED, member of the supervisory board) and the ranking of average willingness to take risks that follows from our vignettes. Our respondents rate the average CEO as more willing to take risks than the average CFO and NED, who are not perceived to be significantly different. In our investment scenarios, the average CFO and CEO ask the same return, while the average NED asks significantly more. The real divide in the boardroom appears to be between the executives and the non-executives. Simultaneously, the average CEO and NED report to experience more risk of a given investment than the CFO. Together these findings indicate it is actually the CFO that is most willing to take risks, followed by the CEO and, after some distance, the NEDs. Since we control for demographics and investment specifics, this effect seems to be driven by the role of the director in the decision process.

Using the survey responses we show corroborating evidence for the use of demographic variables as proxies for the cognitive styles and knowledge base of board members at the level of individual decisions. Our results indicate that age and experience decrease the willingness to take risks of individual board members, which was also found through more indirect methods (e.g. Berger et al. (2013); MacCrimmon and Wehrung (1990); Sanders (2001); Wiersema and Bantel (1992)). We also find that (over-)confidence has a clear negative effect on the return demanded by our respondents and the risks they perceive from investment. This effect is not driven by a specific member of the TMT: we find no significant differences between the distribution of confidence in the groups of CEOs, CFOs and NEDs, and all groups show signs of overconfidence.

The rest of the paper is structured as follows: Section 4.2 discusses some related literature, section 4.3 details our methodology and survey, section 4.4 discusses some empirical predictions that are tested in section 4.5, section 4.6 discusses the results and concludes.

4.2 Related Literature

Most closely related to ours are papers that study the effects of individual characteristics of the board members on their decisions and company outcomes. Many examples can be found in behavioral finance and upper echelon theory (see Barberis and Thaler (2003), Hambrick (2007) and Baker and Wurgler (2011) for surveys). Getting access to the members of the corporate elite is the biggest challenge to this type of research. Leblanc and Schwartz (2007) give an overview of the most used solutions. In this paper we present the results of a first study that combines the possibilities of dynamic webpages with more traditional survey techniques. This methodology could allow a new way of studying the behavior of hard to reach respondents.

Fehr and List (2004) and List and Mason (2011) tackle the problems of access and incentives through a very creative, and effective solution. They visited a conference of Costa Rican coffee companies, which solved the issue of bringing the CEOs together. Due to the wage differences between Costa Rica and the US these authors were able to provide credible monetary incentives to this group of corporate executives, allowing them to run lab experiments. Fehr and List (2004) study the effects of incentives and punishment on the trust and trustworthiness amongst this group of Costa Rican coffee CEOs. They find that these CEOs are more trusting than a reference group of students, but respond similarly to the possibility of using punishment. List and Mason (2011) show that this same group of CEOs presents large deviations from expected utility theory and that their risk aversion seems to increase when the probability of large losses increases. To the best of our knowledge, these are the only papers to present direct laboratory evidence on the risk attitudes of executives.

In a survey study MacCrimmon and Wehrung (1990) find that success is positively correlated with willingness to take risk. They also find that failed risky decisions obstruct career plans. Furthermore, more successful executives rate themselves as more willing to take business risks and rank higher on a sensation seeking scale than the less successful executives.[2] The willingness to take business risk and the need to seek sensation are both negatively correlated to age. We present corroborating evidence for the last finding, older respondents tend to be less willing to take risks in our vignettes.

Graham et al. (2013) administer psychological tests as well as ask their survey respondents about their companies' characteristics to be able to relate the two and guarantee anonymity. Their results indicate that personality traits of the CEO, like overconfidence and optimism, have the expected effect on leverage and other corporate choices. While patience and risk tolerance of the CEO are negatively related to the fraction of flat salary in his total compensation package.

Ben-David and Graham (2013) use statements gathered from CFOs over time to measure CFO overconfidence. Their overconfidence proxy is created by asking CFOs for an 80% confidence interval on the market return at some future moment. Individuals that are too confident in their own prediction will specify too narrow an interval. Of the 13.300 observations 64% of market returns was outside of the provided interval, showing a large amount of overconfidence in the CFO population. The authors show that overconfident CFOs use lower discount rates to value cash flow, invest more, acquire more other firms, use more debt, are less likely to pay dividends and more likely to repurchase shares, and have a higher ratio of long term to short term debt. We use survey-based measures for confidence and relate this to behavior in vignettes.

Instead of surveying the executives, Malmendier and Tate (2005a,b, 2008) use proxies constructed from publicly available data, like CEO stock-option holdings or press statements, to classify CEOs as overconfident. Using these proxies they find that overconfident CEOs overpay for target companies and undertake value-destroying mergers more often, overestimate the returns internally and believe outside investors undervalue their companies. Based on our proxy for overconfidence we find that overconfident board members require a lower expected return and perceive less risk of a given investment, and thus in-

[2] The Sensation Seeking Scale was developed in Zuckerman et al. (1964) and asks respondents to make a choice between two opposing statements, for instance: "I would like to try parachute jumping." and "I would never want to try jumping out of a plane with or without a parachute."

deed take more risks. Our findings indicate, however, that overconfidence is not restricted to the CEO or CFO, but is present throughout the board.

The literature listed above focuses either on the CEO or the CFO. Jensen and Zajac (2004) already argue that to get a full picture of top level decision making, one should study all the individuals involved, which includes both our NED and TMT respondents. Furthermore, we account for the role a respondent has in the boardroom. The importance of this role for individual behavior has received relatively little attention in the literature. The effect is, however, strikingly visible in the experiments of Gillette et al. (2008). Ex ante similar student subjects are randomly assigned a role as supervisory or executive director. The authors show that outside directors can fulfill a watchdog role, and their effectiveness depends on the board structure. Even in this setting with ex ante similar individuals, assigning them a specific role affects behavior. As our findings will show, these role effects are also visible in the behavior of experienced and professional executives. However, the role patterns do not match commonly held beliefs about willingness to take risks. It appears that it is the CFO rather than the CEO that is willing to take most risks in the boardroom. This result is a clear reminder that researchers and practicioners alike need to carefully check their assumptions about individual behavior, so far we have not been able to find (m)any individuals that correctly predicted this relationship.

4.3 Methodology

We collected our data using a survey that ran off a dynamic website. English translations of the questions can be found in Appendix C.1. Our survey ran via a survey website, Erasmus University Consumer and Leading Indicator Database (EUCLID), that is managed on behalf of Erasmus University. This allowed us to approach all potential respondents via their personal e-mail address while guaranteeing their anonymity. Before the data was used for analysis the website administrators removed all indicators that could be used to identify individual participants.[3] Each respondent received a hyperlink with a unique token that allowed the respondents to start, break-off and continue their survey at their own convenience. Before sending the survey to our respondents we created several prototypes of the survey. These prototypes were tested first on people within academia that had some boardroom experience, and later on a small sample of executives enrolled in an educational program at the University. With these individuals we were able to ask

[3]In the interest of full disclosure: one of the authors, Rene Segers, is partner at the company that manages the EUCLID website. To ensure the integrity of the EUCLID website, his firm and this research project, the identifying tokens and the email addresses were at no point part of the created research database.

post survey questions to find out whether the manipulation of the investment sizes used and the way of presenting the questions had the desired effects. Only after their responses confirmed that the questions were clear and that the vignettes were realistic enough, we proceeded to send the survey to our respondents.

The token in the hyperlink allowed us to identify individual respondents, randomize the order of the vignettes, and use the possibility of dynamic web design to tailor the vignettes to every single respondent. Unlike standard survey data about risky investments, our measures do not suffer from recollection biases, or biased reporting. By using the same method to calculate initial investments and presenting the investments in the same way to all subjects, our approach makes it possible to compare realistically sized investment scenarios both within and between subjects. Simultaneously, illusion of control, a high degree of commitment to good outcomes, and abstract or unknown reference points make it hard to compare performance across individual in a traditional survey even without measurement errors. Since recollection bias does not play a role and the variation is exogenously determined by the researchers, these biases do not occur in our survey.

Even though the questionnaire was administered by an affiliated third party (which was clearly visible on the webpage), the emails were all sent in the name of the first author, a professor at Erasmus University. Since many of the email addresses used were gathered via the network of the University, mostly through alumni networks and business networks connected to the city of Rotterdam and the University, attaching the name and logo of the Erasmus University gave the emails some extra credibility.

We cannot incentivize the respondents through monetary earnings, as is usually done with student subjects and as was done by Fehr and List (2004) and List and Mason (2011) with executives. The simple reason is that a substantial part of our respondents earn more than our research budget, such that any money we offered would likely have been considered an insult. This constraint will likely remain if one wants to do research on Western executives. Instead of a monetary reward, we offered our respondents an individualized report of the results of the research they were involved in. These reports showed the respondents both the distribution of the answers given and their position within this distribution for the relevant questions, allowing them to learn more about themselves in relation to their peers. The need to collect sufficient responses to generate the repports, created some time between the survey and the feedback. This also ensured that the easiest way for our respondents to reflect on the results of the survey, was by filling in the survey honestly. Furthermore, because of the external management of the database and guaranteed anonymity, there was no strategic reason to lie. Neither the researchers nor any other respondent can find out how any particular respondent replied, such that answering honestly cannot have direct negative consequences. We believe that the relatively high response rate, a gross response rate of 43% (see subsection 4.3.2) compared to a response rate of roughly 10% in other recent surveys, (e.g Graham et al. (2013, 2005); Trahan and Gitman (1995)), shows that the ability to learn about themselves and their environment is an important reason for many of our respondents to reply, and reply honestly.

4.3.1 Survey

The survey first elicits some basic information from the respondents (questions 1 to 4). We ask for their position within the company, the approximate size of the company revenues (with five choice options, ranging from below €50 million to more than €1 billion), the company sector and number of employees (ranging from less than 50 to more than 10.000). We then ask the respondents to indicate the number and size of investments they have decided upon in five categories (new market expansion, expansion of production capacity, R&D projects, ICT projects, and mergers and acquisitions) in the last 15 years (questions 5 and 6).

Respondents who indicate to have experience with investments in a particular category are asked to evaluate two vignettes with investment possibilities in this setting (except for M&A). The investment possibilities are described as originating from the relevant department in their company and these hypothetical investments are paired with an estimated probability success (0.80 or 0.95) and a non-recoverable investment cost. The size of the investment cost depends on the typical size of investment decisions made by the respondents (derived from the first part of the survey) multiplied with a random factor (0.5 or 1.25). Respondents are asked to indicate the minimum Net Present Value (NPV) they, in their professional role, would require to approve of this investment (questions 7 to 14). After they have done so, they are asked to self-reflect and indicate "how risky" they perceive this investment to be on a 7-point scale.

Based on Dohmen et al. (2011) we ask our respondents to rate, on 11-point scales ranging from "Not willing to take any risk" (0) to "Very willing to take risks" (10), how they perceive the willingness to take risks of: themselves in general, of themselves in their professional role (where the role they have is inserted automatically by the survey), of the average CEO, of the average CFO, and of the average NED (questions 15 till 19).

At the end of the survey respondents are asked about their general background, gender and age. A final question asks whether respondents used a calculator during the survey.

4.3.2 Respondents

We gathered a panel of CEOs, CFOs, NEDs and a reference group consisting of (non-executive) managers, consultants and analysts of roughly the same size(appendix C.2 shows the proportions of roles in our survey). The individuals in our survey work for Netherlands based national, international or multinational companies, both listed and unlisted, in a range of industries. All individuals reported experience in investment decisions for their company. Our database contains personal data and it is protected. That is, the listed persons are treated anonymously and the database is for scientific use only.

We sent the questionnaire to 682 respondents of whom 43% responded. 160 (23%) Respondents completed the entire survey and it took them on average 17 minutes to do so. Another 131 (20%) respondents did not answer all the questions (they spent on average 9 minutes on the questions), but we can still use their answers for part of the analysis. Summary statistics can be found in Appendix B. To reduce potential noise in the responses, we removed respondents who:

- demand the same NPV or report the same risk perception on all vignettes, demand an NPV above €100.000 mln, or demanded an €0 NPV for one or more of the investments choices (n=25);
- based on their reported birth date are minors (n=1);
- have a confidence proxy equal to - 6, since this was a clear outlier (n=1).

All of the reported results are based on the remaining data.

4.4 Empirical predictions

Our vignettes are meant to simulate realistic investment scenario's. If our respondents treat the vignettes as such, both the NPV and the perceived risk should increase if the probability of success decreases and if the investment size increases.

Prediction 4.1. *The NPV and perceived risk increases in both the initial investment and in the probability of failure.*

Depending on the experience of the participants, we asked the participants to consider investments in R&D, market expansion, IT, and/or an increase in production capacity. Some authors argue that different types of investments are treated differently by the corporate executives. Sanders (2001) indicate this could be caused by option based compensation pay. This pay structure yields an incentive to place large bets on risky type of investments. If the investment categories have different risks, they should be differently affected. However, differences in investment behavior could also be due to experienced differences between the types of investment. If risk attitudes are influenced by the context of the investment as suggested a.o. by Malmendier and Tate (2008), it would be virtually impossible to disentangle these effects in observed behavior alone. As we do not use (option based) incentives, the observed differences can only be due to the context. In our setting there is no a priori indication about which investment types would be considered more risky, while all respondents only get vignettes in investment types they report to have experience with, so we predict that no difference is made between the different categories:

Prediction 4.2. *Neither the NPV nor the perceived risk is influenced by the investment type.*

The question about general willingness to take risk, was validated in Dohmen et al. (2011). Based on their results, individuals that rate themselves higher on this scale, have a higher risk tolerance.

Prediction 4.3. *The NPV and perceived risk decreases in the risk tolerance of individuals.*

In the upper echelon theory and behavioral finance literature, indirect measures and demographic characteristics of individual executives are used to explain corporate outcomes like stock behavior and merge decisions. This approach treats the company as a black box, which makes interpreting the uncovered relationships challenging (Lawrence (1997); Huse (2005)). Because of the potential confounding effects, one has to be careful when using individual characteristics as proxies to explain company outcomes. Our survey provides a complementary approach to this type of research by probing the black box of the company. We can test some of the effects identified in this literature directly on the individual decision makers.

Based on earlier findings we predict that the willingness to take risks decreases in age and experience, such that the NPV and perceived risk should increase in both age and experience (Berger et al., 2013; MacCrimmon and Wehrung, 1990; Sanders, 2001; Wiersema and Bantel, 1992).

Prediction 4.4.

a The minimally required NPV and perceived risk increase in age.

b The minimally required NPV and perceived risk increase in experience.

The first 4 predictions treat all respondents similarly, however, one of the biggest advantages of using dynamic web surveys is that it allows us to present respondents with investment scenarios that differ exogenously, while still tailoring each scenario to the experience of our respondents. Since every respondent's experience, background, and risk propensity are part of why he has a certain professional role, the decisions taken by each respondent in their professional careers are at least in part endogenous. The variance in investment size and success rates in the vignettes is exogenously imposed by the experimenters, which allows us to get a cleaner measure of the effect of professional roles on individual decision making under risk.

There is relatively little other evidence on the effect of professional roles on behavior of executives, while there is also no consensus in theory. It is often assumed that the CEO is the Entrepreneur on the board, while the NEDs are supposed to act as 'some sort of discipline' (Mace, 1972), to a lesser extend this later role is supported by the CFO as 'bean counter' (Favaro, 2001; Zorn, 2004). This division of tasks calls for more cautious NEDs and CFOs and thus increases the demanded NPV from agents in those roles. Conversely, in much of the agency theory literature that deals with the remuneration contract of the CEO/manager, it is assumed that risk neutral investors (who are represented in the remuneration boards by the NEDs) contract with risk averse CEOs (Shavell, 1979; Stiglitz, 1987a; Holmstrom and Milgrom, 1987). The existence of option contracts could then be seen as indications that the risk neutral NEDs are trying to incentivize a risk averse manager. This would indicate the opposite conclusion, namely that the CEO (and to lesser extend) and CFO are more risk averse than the NEDs. Without clear theoretical guidance, we rather skip ahead to the rankings given by our respondents, which we will discuss just below. Based on the survey responses, we predict that CEOs are most risk tolerant, while NEDs are least risk tolerant, with CFOs somewhere in between.

Prediction 4.5. *Both the NPV and the perceived risk is larger for NEDs than for CFOs, and for CFOs than for CEOs.*

4.5 Analysis and results

We first analyze our respondents' rankings of self-reported and perceived willingness to take risks. The average risk tolerance scores are depicted in Figure 4.1, where we distinguish between the answers given by CEOs, CFOs and NEDs. Significant differences (based on t-tests) are shown through the arrows. The dotted arrows represent a significant different at a 95% level, while a solid arrow denotes a significant difference at the 99% level. In their professional roles CEOs report a higher willingness to take risks than

CFOs and NEDs, as can be seen in Figure 4.1a. We do not find significant differences between the self-reported risk tolerance of CFOs and NEDs. The reported willingness to take risks in general, depicted in Figure 4.1b, yields the same conclusions. CEOs again report a higher willingness to take risk than CFOs and NEDs. We do not find a significant difference between the answers given by CFOs and NEDs.

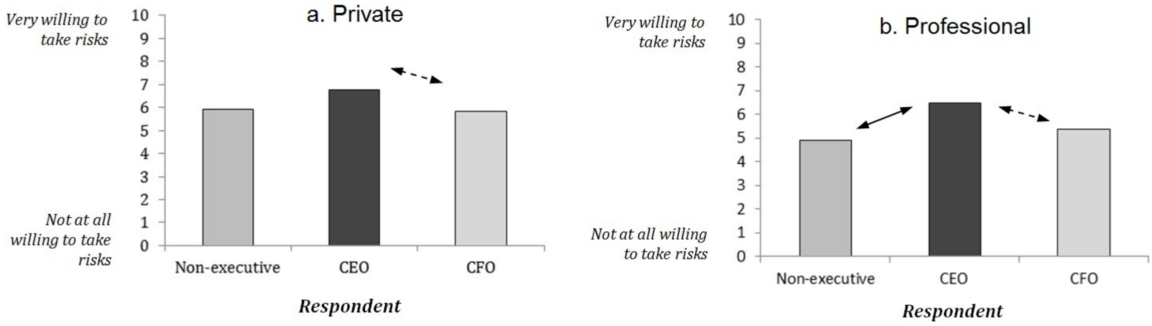

Figure 4.1: self-reported willingness to take risk

Figure 4.2 shows how the board members perceive each other's willingness to take risks (questions 17 to 19). The results are in line with the self-perceptions in Figure 4.1. CEOs are considered to be significantly more willing to take risks by all members of the corporate elite. CFOs and NEDs are perceived to have more or less the same risk tolerance on average.

4.5.1 vignettes

In the analysis of the investment scenarios, or vignettes, we examine whether the perceived differences in risk tolerance across the three roles of the corporate elite are associated with the reported investment choices. We regress the natural logarithm of the minimally required NPV (model (1)) and the experienced risk (model (2)) on a variety of possible explanatory variables. For both models we consider the size of the investment, the probability of success and the investment category as investment specific variables and we include individual characteristics such as age, the respondent's professional role, and a set of measures of the respondent's experience level. Both models are estimated using unbalanced panel techniques including individual random effects.[4]

[4]The parameters of model (1) can be estimated relatively straightforward using maximum likelihood, see Breusch (1987) among others. The perceived risk is rated on a multinomial scale. Model (2) belongs to the class of generalized linear mixed models (GLMMs). The parameters of this model can conveniently be estimated using the Penalized Quasi-likelihood method Rabe-Hesketh et al. (2004, 2005). Estimation methods for our two models are available in Stata through the packages xtreg and gllamm, respectively.

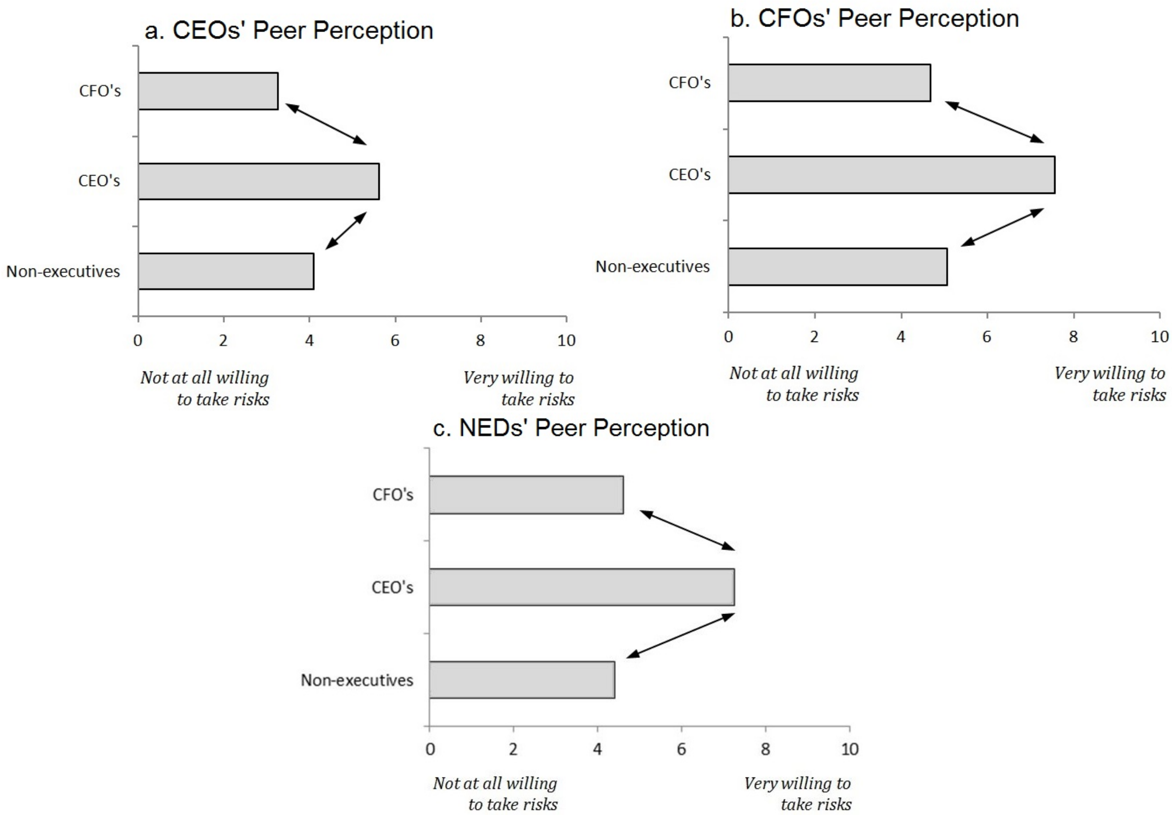

Figure 4.2: Perceived willingness to take risk

The estimated coefficients of model (1) and (2) (excluding auxiliary parameters and random effects), together with their standard errors, are presented in Table 4.1. For identification purposes we use ICT-projects as the reference investment category. As the reference group for the individuals we consider the pool of analysts, managers and consultants (i.e. respondents outside of the board).

The demanded NPV has an estimated elasticity of approximately 0.7, implying that a 1% increase in the size of the initial investment translates to a 0.7% increase in the demanded NPV. Respondents report to perceive more risk when they are confronted with larger investments. As the probability of a successful investment outcome increases, individuals are prepared to accept a lower NPV and experiences less risk. In both models, these two effects are significant at the 1% level. Thus, we find support for predictions 4.1, showing that our manipulations had the expected effect.

In line with prediction 4.2, the differences between the four types of investment are not significant (Wald test p-value ≥ 0.16). The table does not present the results for Prediction 3, since the coefficient on the risk propensity measure is consistently insignificant. Although Dohmen et al. (2011) spent considerable effort in validating the question,

Variables	1 Log(NPV)	2 Perceived risk of investing
Properties of the investment		
Log(Investment)	0.698***	0.234***
	(0.0391)	(0.0387)
Probability of success	-1.818***	-5.875***
	(0.493)	(0.555)
Type of investment: expansion to a new market	0.0694	0.00572
	(0.106)	(0.112)
Type of investment: R&D-project	0.0366	-0.0359
	(0.103)	(0.110)
Type of investment: increase in capacity	-0.0758	-0.0324
	(0.104)	(0.109)
Individual characteristics		
Role: non-executive	0.387	0.892***
	(0.418)	(0.237)
Role: CEO	-0.310	0.714**
	(0.348)	(0.280)
Role: CFO	-0.703**	-0.0808
	(0.308)	(0.180)
Age	0.272**	0.288***
	(0.116)	(0.0730)
Age^2	-0.00306**	-0.00284***
	(0.00122)	(0.000762)
Avg. size of previous investments	0.0446**	-0.0338*
	(0.0220)	(0.0173)
Avg. size of previous investments^2	-0.000496*	0.000445**
	(0.000285)	(0.000208)
Avg. number of previous investments	0.0905	-0.0532
	(0.0937)	(0.0605)
Used a calculator	-0.858***	-0.342**
	(0.315)	(0.175)
Number of observations	716	758
Number of individuals	113	117

The superscripts ***, **, * denote significance at the 1%. 5% and 10% level, respectively. Standard errors are in parentheses.

Table 4.1: Estimation results of models (1) and (2)

it seems its predictive value in our smaller sample is negligible. The measure did contain some predictive value after a transformation, however this changes the interpretation of the measure considerably. We discuss this new measure in the next subsection.

After controlling for the role effect we still find a significant effect of age in the data. As professionals get older they tend to become 'sadder but wiser' in the sense that they become more careful investors. The demanded NPV increases as well as the perceived risk. This effect diminishes over time and it peaks at the age of 45 for the NPV and 52 for the perceived risk, as indicated by the negative quadratic terms. In line with the age effect, we find that investors who have experience with higher investments levels tend to be more cautious by demanding a higher NPV. However, the opposite is true for the level of perceived risk. Investors who have more experience with higher levels of investments perceive less risk. There seems to be no additional effect of experience measured in terms of the number of investments. Thus, we find partial support for predictions 4.4a and 4.4b and our findings indicate that using observable demographics as proxies for individual behavioral characteristics, as is done in behavioral finance and upper echelon theory, can be a viable approach.

The most novel result of our survey is the possibility to identify the effect of professional roles. Looking at table 1 it becomes clear that NEDs ask a higher NPV than CEOs (Wald test p-value is 0.09) and CFOs (Wald test p-value is 0.01) and thus demand the highest risk premium overall. Based on our survey results, we hypothesized CFOs to be more risk averse than CEOs, but both groups tend to demand the same NPV (Wald test p-value is 0.22). In terms of NPV there seems to be a sharp distinction between the NEDs on the one hand and the executive board members on the other hand. In contrast, when we look at the perceived risk metric, the divide is between the risk tolerant CFOs and the relatively risk averse CEOs and the NEDs. The responses to our vignettes thus do not support prediction 4.5. Like our respondents expected, NEDs are consistently most risk averse, both in terms of demanded NPV and in terms of perceived risk. Unlike our expectations, however, CFOs are more risk tolerant in both dimensions than NEDs, and even weakly more risk tolerant than our CEOs.

Finally, respondents who reported the use of a calculator tend to ask a lower NPV and perceive less risk. This can plausibly be contributed to a feeling of control. Since these respondents could calculate the expected return on investment and, as a result, felt more in control, they might have perceiving less risk.

4.5.2 The effect of (over-)confidence

The responses to the risk tolerance question (question 16) do not predict individual behavior very well. However, for our executive respondents we can transform the measure. We subtract the risk tolerance rating the board members gave to their average peer from the rating they gave themselves (question16 - question17/18/19). Both ratings are on the same scale, such that we could take the simple difference to normalize the measure. The transformed measure ranges in value from -2 till 5 (after deletion of the single -6 observation), with a mean of 0.78 and a mode of 1, figure 4.3 shows the entire distribution of our confidence measure.

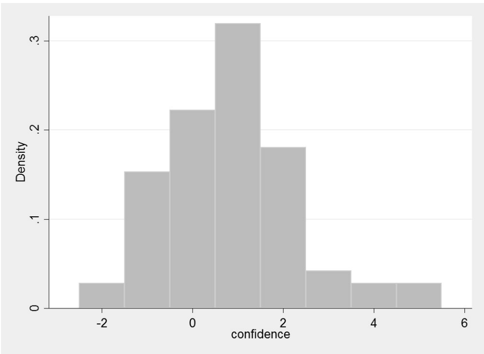

Figure 4.3: The distribution of confidence in the board

Through this normalization the interpretation of the measure changes. A CEO that claims to be more willing to take risks than the average CEO, is not just saying he is willing to take risks. He compares his risk tolerance to that of people in the same position, for all of whom making risky investment decisions is part of the job. Excessive willingness to take risks can be explained in several ways. It could just indicate that the executive is simply less risk averse than the average person, however, since our measure has a mean and mode above 0, this would create a statistical anomaly. It is more likely, therefore, that this measure picks up on a form of (over-)confidence. Either our subjects overestimate their own ability or risk tolerance because they show 'peer group neglect' (Camerer and Lovallo, 1999), or it could simply measures a better-than-average type of overconfidence. We therefore interpret this measure as a proxy for overconfidence.

Overconfidence is an important driver of individual choice behavior (Griffin and Tversky, 1992), in earlier studies it has been studied in CEOs and CFOs. However, based on a Kolmogorov-Smirnov test we do not find significant differences across the groups of board members, including NEDs. To assess whether an individual's confidence level drives the demanded NPV and the perceived risk in the investment scenario's, we included our confidence metric in models (1) and (2) specified above, and reran our regressions. In these regressions, we treat our pool of CFOs as the reference group. The results are presented in Table 4.2.

Before we turn to the interpretation of the results, note that our samples are now approximately 30% smaller as overconfidence is only measured for board members, such that we have to drop our reference group. This causes the standard errors in the regressions to be larger than in Table 4.1.

In this second set of regressions we again find a strong effect of the size of the outlay and the probability of success on the demanded NPV and perceived risk. The differences between the professional roles are also still apparent from the results. While the effect

Variables	1' Log(NPV)	2' Perceived risk of investing
Properties of the investment		
Log(Investment)	0.662***	0.184***
	(0.0458)	(0.0450)
Probability of success	-1.944***	-5.876***
	(0.609)	(0.650)
Type of investment: expansion to a new market	0.112	-0.00409
	(0.131)	(0.131)
Type of investment: R&D-project	0.0740	0.0395
	(0.128)	(0.129)
Type of investment: increase in capacity	-0.0238	0.0615
	(0.128)	(0.128)
Individual characteristics		
Role: non-executive	0.761**	0.854***
	(0.388)	(0.270)
Role: CEO	0.249	0.628***
	(0.294)	(0.242)
Role: CFO		
Age	0.0389	0.226*
	(0.154)	(0.120)
Age^2	-0.000642	-0.00226*
	(0.00160)	(0.00120)
Avg. size of previous investments	0.0648***	-0.00450
	(0.0234)	(0.0180)
Avg. size of previous $investments^2$	-0.00074**	5.22e-05
	(0.000297)	(0.000220)
Avg. number of previous investments	0.0366	-0.118
	(0.102)	(0.0759)
Used a calculator	-0.914***	-0.212
	(0.352)	(0.216)
Confidence	-0.240***	-0.186***
	(0.0861)	(0.0704)
Number of observations	526	552
Number of individuals	79	81

The superscripts ***, **, * denote significance at the 1%. 5% and 10% level, respectively.
Standard errors are in parentheses.

Table 4.2: Estimation results of models (1') and (2') with the confidence metric included

of age and the usage of a calculator are weakly significant in model (2'), the effect of experience, as measured by the typical size of previous investments, only survives in model (1'). This is probably due to the smaller sample size. All of our main conclusions seem robust to inclusion of measures of individual confidence.

Finally, and most importantly, we indeed find a strong and negative effect of confidence both on the demanded NPV and on perceived risk. Board members that are more confident about themselves tend to demand a lower NPV, thus requiring a smaller risk premium for their investments. Additionally, they perceive less risk than their less confident peers.

4.6 Conclusions

In this paper we studied the risk perceptions and risk assessment of CEOs, CFOs and non-executives using a dynamic web based survey. We found that CEOs are generally perceived to be more risk tolerant than CFOs and NEDs, however, when we asked our respondents to assess realistic investment scenarios tailored to their personal experience, the picture changed. Firstly, non-executives demand more return on investment than executives. Comparing the CEO to the CFO, we find that CEOs report to perceive more risk than CFOs, and demand non-significantly more returns, indicating the CFO is most willing to take risks. These results remain after controlling for individual and investment characteristics.

This discrepancy between peer perception and individual risk assessment could lead to miscommunication in the boardroom. It suggests that practitioners as well as researchers should be careful about what we assume about the relative risk tolerance of board members. The relationship might not be what we expect and can be influenced by professional roles.

Role dependent risk attitudes can have several causes. Selection effects, leading individuals with similar risk tolerance to choose the same role in board could play a role (Goel and Thakor, 2008). It seems unlikely, however, that this explains our findings, as we control for individual characteristics. Moreover, all directors operate in a high stakes environment and many of the NEDs have TMT experience. A more likely explanation is the adaption of professional roles by directors. If we look at the role of the NED as a safeguard, he or she might feel the urge to be more careful than the TMT. As NEDs can be assumed to be aware of their role in the decision process, they will likely act accordingly. Information asymmetry is another, closely related, explanation (Lorsch and Young, 1990; Lorsch, 1995). If NEDs know that in general they do not have the same amount of information as the TMT, one response to this problem is to create extra (financial) safeguards. This means that when a board member is approached in his or her

NED mindset, he or she starts creating these safeguards as a matter of course, and a more cautious professional role is created as a strategic response. Especially for NEDs who serve on several subcommittees, like remuneration and audit committees, the adoption of different professional roles could be a good mechanism to cope with the different requirements posed by the different tasks.

Our findings support the basic premise of behavioral finance and upper echelon theory that demographics can be used as proxies for individuals' cognitive styles and knowledge base. We look at individual choice behavior of corporate directors and find that age and experience lower the willingness to take risk, while (over-)confidence increases risk tolerance, as was found through more indirect means in earlier papers.

4.6.1 Limitations and future research

When studying the upper echelons of the corporate sector it is challenging to bring together a large group of relevant respondents. By using a dynamic website we are able to apply several methodological features of lab research to a group of respondents that is almost impossible to get into the lab. The dynamic website allows us to tailor each survey to the respondent, while guaranteeing anonymity, and randomizing treatments.

Although it is virtually impossible to incentivize high ranking corporate officials through monetary payments, as is usually done in experimental economics, we believe the individual reports we generated provide a good reason to reply honestly. The possibility to learn about oneself and ones' peers is probably worth more to many key decision makers than the small sums of money usually offered to experimental subjects. The fact that, with our limited number of respondents, we are able to find corroborating evidence for earlier findings from the behavioral finance and upper echelon literature, shows that indeed our survey measures the effects we expect to be there. Simultaneously, the exogenously imposed variation allows us to compare both within and between respondents, allowing us to answer more types of questions than with normal surveys.

Like any methodology this one has its weaknesses. We do, for instance, not have the same amount of control as lab experiments, nor can we easily study interactions, or easily rerun a failed 'treatment'. This means researchers have to do without some of the advantages and methods of experimental economics. We deal with this problem by extensively testing prototypes of our survey, first on people within academia and then on a small group of board members from an educational program on the Erasmus (i.e. in a controlled setting were we could question the respondents afterwards). Because of the dynamic nature of the survey, the questions change between subjects. This increases the time and effort required for developing the survey considerably. A stronger test of the strength of the methodology would be to link the findings to actual, rather than hypothetical, choices. This is quite difficult in our setting due to both the limited number of observations and the anonymity we guarantee.

An issue we failed to account for in our prototype testing was the time aspect. Our respondents could fill in the survey at their own leisure and the survey software allowed them to break off and complete the survey later, without losing their responses up till that point. Despite the convenience, our attrition rate is close to 50%. In retrospect we believe that we should have paid more attention to this aspect. We severely underestimated the time it took a respondent who was unfamiliar with the survey environment to complete the entire survey, and this might have caused a higher than necessary attrition rate. This issue had received some, but apparently not enough attention in designing the survey. During the prototype tests neither the academic respondents, who were quite well aware of the importance of the test, nor the executives and managers, for whom it was part of their homework for the course, seemed to have minded the time it took. Clearly, our actual respondents felt differently.

Lastly, it is hard to create direct interaction between subjects/respondents through a website. Especially with executives it seems unlikely that one can get a significant number of respondents to log on at the same time. An alternative would be to develop games that are played sequentially. Then each respondent would only have to continue when the other player(s) has (have) made his (their) move(s). For this first pass, however, we focus on decision making under risk, an area of research where this aspect is less relevant, but alternatives exist and this could be an interesting avenue for future research.

In this paper we found that the behavioral tendencies like overconfidence and the effect of age and experience are prevalent throughout the board. Existing research has focused primarily on the CEO. We believe that future research should try to focus on all members of the board. We hope therefore that the fact that we find signs of overconfidence for these other board members encourages researchers to start this with this hopefully fruitful line of enquiry.

Chapter 5

Can (Managerial) labor Markets undo Signal Jamming by Career Concerned Managers? Experimental Evidence[1]

5.1 Introduction

The idea that managers strive to "provide the most positive signals to the managerial labor market" to influence this market's perceptions of their abilities and that "markets appropriately use current and past information to revise future wages" was one Fama's (1980) key contributions to the debate on how the separation of ownership and control can be an efficient form of organization.[2]

Holmström (1999) was the first to explore this idea formally.[3] His analysis shows that in equilibrium the behavior of the manager and the market fit together in a subtle way. With rational markets, any efforts of the manager to inflate the market's perception are to no avail. The labor market rationally anticipates that the manager uses his actions to influence any signals the market can somehow relate to the manager's ability. Nevertheless, although the manager cannot fool the market, this does not stop him from trying to impress the market. Not doing so would damage the inferences that the market draws from the signals it receives, and thus damage the manager's career.

[1] This chapter is based on joint work with Bauke Visser

[2] The quotes are on pages 293 and 296, resp. See also "In short, (...) we impute efficiency or rationality in information processing to managerial labor markets" on page 296.

[3] Holmström's piece was originally published in a hard-to-find Festschrift in 1982.

On the theoretical front, this idea has proven very productive. In his seminal contribution, Holmström showed how market-based incentives or career concerns may lead the manager to exert additional effort (as future marginal benefits complement current ones), but can also give rise to distorted investment or project choices. Subsequently, models have been developed that explain e.g., how the nature of the ability that is important for a manager's career affects the way in which career concerns influence a single manager's actions,[4] how career concerns give rise to herd behavior and the negligence of private information among groups of managers or financial analysts,[5] and how such concerns obstruct the exchange of information in decision-making committees, influence the votes cast and lead such committees to show a united front to the outside world.[6]

Despite the importance of the equilibrium mechanism uncovered by Holmström and others – managers jamming signals and markets undoing this jamming – very little is known about its existence and operation in reality. The main reason is clear: to observe whether managers distort decisions to influence market perceptions, one should know what the correct decision is. This is typically hard to establish for an outsider. Probably as a result, empirical work has limited attention to intertemporal patterns of a manager's compensation that can be explained by career concerns.[7] As uncertainty about a manager's ability declines over time, the sensitivity of market-based compensation to observed performance should go down as well. To compensate, a firm should offer the manager incentives that get stronger with his tenure. Gibbons and Murphy (1992) find evidence for this pattern in a sample of executives of US-based companies. Chung et al. (2012) and Lim et al. (2013) find evidence that market-based rewards of private equity and hedge funds managers also decline with tenure. In the case of financial analysts, the information asymmetry about what the right decision was at the the time of the decision can be reduced ex post. As there is always a number of analysts looking at the same stocks, one can use the actions of the other analysts to get an idea of what the right decision is and when distortions occur. A very influential contribution in this line of research is Chevalier and Ellison (1999), who study the termination decision of mutual fund managers and show that younger managers both have and heed an incentive to herd into popular sectors. This paper is followed, a.o. by Hong et al. (2000), who show similar herding incentives and behavior with respect to securities analysts earnings forecasts. Clement and Senyo (2005)

[4]For example, in Milbourn et al. (2001) managers differ in their ability to develop profitable projects; in Suurmond et al. (2004) managers differ in their ability to identify the state of the world.

[5]See e.g. Scharfstein and Stein (1990) and Ottaviani and Sørensen (2001) for models in which agents are privately informed at zero cost.

[6]See Visser and Swank (2007).

[7]For a discussion of confounding factors in this line of research, see Hermalin and Weisbach (2014).

extend this line of research to include analysts' characteristics. These authors show that more experienced, more precise, more specialized analysts are more likely to deviate from the herd and that deviators are more precise on average.

In this paper, we present the results of a laboratory experiment that was designed to study Fama's idea and the mechanism identified by Holmström and others. As in Fama (1980), we study a management committee rather than a manager operating in isolation. The set-up closely follows the model of Visser and Swank (2007). In the model, a committee has to make an investment decision. Each manager receives a private signal about the profitability of the investment. During the meeting, they discuss whether to invest or not, and next cast a vote. Managers care about the profitability of the investment and the way the decision affects their careers. These career concerns are summarized by the updated belief held by the labor market that a manager is highly able at assessing the quality of investment projects. Once the decision has been taken and observed by the labor market, a manager can send a statement to the labor market about the degree of confidence he has in the decision, how they reached the decision etc. This means that, just as in Fama's account, labor markets receive many signals that they have to sift through to determine the ability of the managers. What the market does not observe is whether the investment actually pays off or not. This is consistent with Fama's focus on top management and the fact that the consequences of decisions made by top management often do not become visible until after quite a while. Thus, markets are facing a hard inference problem. Nevertheless, the model provides clear predictions that can be tested in the lab.

In the experiment, half of the subjects are managers who decide in pairs whether or not to invest. The other half of the subjects are market participants, who bid for the services of the managers after observing the decisions of the managers and the statements they make. We find convincing evidence that managers want to fool the market and that market participants take this into account in their bids. In line with theory, bidders value managers more when they have invested, than when they did not invest. As a result, managers sometimes deviate from maximizing investment profitability by choosing to invest although they know that the expected profit is negative. Market participants see through these attempts to fool them, and the difference in wage offers is close to what theory predicts it should be in the presence of signal jamming (and is indeed smaller than what it would have been without such attempts).

In the experiment, markets also reward managers if managers show a united front by expressing the same high level of confidence in the decision they took. In theory, however, these statements should be considered cheap talk and not lead to changes in valuations:

once the investment decision has been taken, rational managers could costlessly exploit any change in valuation stemming from a change in statement. Thus, in equilibrium, the market's valuation is determined by the investment decision – a costly signal – and not by the cheap talk statements used. Are the market participants fooled by these statements? Or is there useful information about a manager's ability in his statements? Analysis of the statements shows that the latter is indeed the case. Moreover, our analysis suggests that bidders are close to making efficient use of this information.

There are few papers that study experimentally whether reputational concerns affect decision-making. Koch et al. (2009) implement a simple version of Holmström's (1999) effort model: performance of a subject equals the sum of his chosen level of costly effort and her ability level. They compare two treatments, one in which performance and ability are separately observed before "firms" make a wage offer and a second treatment in which firms only observe performance. In line with the theoretical prediction, they find that subjects exert more effort to influence a firm's inference concerning their ability if ability is not observed by firms. Subjects' incentives to jam signals about their ability is, however, only taken into account to a limited extent by firms.

The aim of the experiment of Berg et al. (2009) is to explain the well-established phenomenon of a decision maker's "unwarranted commitment to past choices" (p. 135), i.e. her commitment to a previously made choice in the face of new information that shows that another choice is better. They compare three hypotheses that have been proposed in the literature: internal justification (a need to justify past decisions to one-self), external justification (a need to do so to others or to appear consistent), and reputational concerns (a desire to establish a reputation for ability). By manipulating the informational setting within which subjects make decisions, they find evidence that the third explanation is the better one: subjects stick to unwarranted decisions only if they are privately informed about the state and if there is a market that observes them. We differ from both studies in our focus on committee decision-making and on a multi-dimensional information environment that combines both costly signals and cheap talk signals from which observers have to infer relevant information.

Fehrler and Hughes (2014) is the only other experimental study of reputational concerns in a committee setting we are aware of. Its focus is quite different. They compare how secret and transparent decision making affects the quality of information exchange within the committee, members' willingness to change views in the meeting, and, as a result, the quality of decisions taken. They do so in a situation in which a committee member is only motivated by her reputation and is privately informed about her ability level. Moreover, the principal learns the state before he updates his belief about the ability of committee members. If members deliberate and vote in secret, all private information (signals and ability levels) is pooled and everybody votes in favor of the decision that is most likely to match the state given the available information. Transparency hampers information aggregation in the committee, for example as members don't want to state, or lead the principal to infer, they are of low ability. The experimental results show broad support for theoretical predictions. However, and just as in our experiment, subjects reveal more information about their ability than theory predicts. It is indeed a recurrent finding in the experimental literature involving cheap talk, that subjects are more truthful than they should be, see e.g. Cai and Wang (2006) and Goeree and Yariv (2011).

This chapter is built up as follows. In the next section we describe some theoretical predictions. In section 5.3 we describe the experimental design and the procedures used. The results are discussed in section 5.4. We report on robustness checks in section 5.5. Section 5.6 concludes.

5.2 The Model

This section provides the outline of a simplified version of the theoretical model in Visser and Swank (2007; VS from now on) on which our experiment is based.[8] A two-member committee, with members $i = \{1, 2\}$, has to decide whether to implement a project, $I = 1$, or to maintain the status quo, $I = 0$. By normalization, status quo delivers a project payoff equal to zero. Project payoff in case of implementation is uncertain and state dependent. It equals $p + \mu$, where $\mu \in \{-h, h\}$ with $\Pr(\mu = h) = 1/2$. With slight abuse of notation, we let μ represent both the state and the state dependent value. The parameter p is the ex ante expected payoff from $I = 1$. As $p < 0$, the committee has a bias against project implementation. Therefore, VS call the decision to implement the unconventional decision. Also, $p + h > 0$, implying that the optimal decision depends on the state.

Figure 5.1 depicts the time line of the game. Nature determines both the state, μ, and the ability level of each member, $a_i \in \{\underline{a}, \bar{a}\}$, with $\Pr(a_i = \bar{a}) = \pi$. Where $\bar{a}$ stands for high ability, while $\underline{a}$ denotes low ability. Nature does not inform anyone about either μ or a_i. Next, the game proceeds with four consecutive stages, an information stage, a deliberation stage, a voting stage, and a statement stage. In the *information stage*, each member receives a private signal $s_i \in S = \left\{s^b, s^g\right\}$ about the state μ (b is bad and g is good). A signal refers to a member's assessment, forecast or view of μ. The quality of

[8] See Visser and Swank (2007) for generalizations of many aspects of the model outlined here and for proofs of the results. Other generalizations can be found in Swank and Visser (2013)

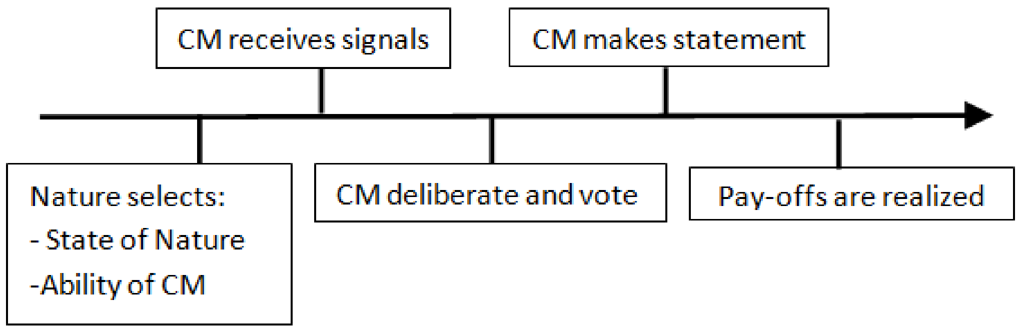

Figure 5.1: Time line, CM stands for committee member.

the signal member i receives depends on his ability, a_i. If i is of high ability, he always receives an informative signal about μ, $\Pr(s^g \mid \mu = h, \bar{a}) = \Pr\left(s^b \mid \mu = -h, \bar{a}\right) = 1$. If i is of low ability, he always receives an uninformative signal about μ, $\Pr(s^g \mid \mu = h, \underline{a}) = \Pr\left(s^b \mid \mu = -h, \underline{a}\right) = 1/2$. In the *deliberation stage*, both members send a message $m_i \in M = \left\{m^b, m^g\right\}$ to the other member. This message may be related to his private signal. In the *voting stage*, members cast a vote on the project, $v_i \in V = \left\{v^b, v^g\right\}$, where $v_i = v^b$ ($v_i = v^g$) denotes that i votes for $I = 0$ ($I = 1$). Of course, i's vote can be based on his private signal and on the messages received. The voting rule in the committee is unanimity, such that implementation requires that both members vote in favor of implementation. The decision taken by the committee becomes observable to the market. Finally, in the *statement stage*, committee members can decide what statement to make to the market concerning what happened inside the committee. This statement could theoretically involve claims about their private signals, the messages they exchanged and the votes they cast – anything to which the outside world would otherwise not have direct access.

Each member cares both about the value of the project and about his reputation (a short hand for career concerns). The reputation of member i is defined as the posterior belief $\hat{\pi}_i(I, \omega) = \Pr(a_i = \bar{a} | I, \omega)$ held by the market that i is highly able. This reputation is based on the investment decision taken, I, and on whatever else the members decide to tell the market, ω. As the market does not observe the state μ, it cannot base its belief on the state. Member i's preferences are represented by:

$$\begin{aligned} U_i(I=1) &= p + \mu + \lambda \hat{\pi}_i(I=1, \omega) \\ U_i(I=0) &= \lambda \hat{\pi}_i(I=0, \omega). \end{aligned} \tag{5.1}$$

The parameter λ measures how much a committee member cares about his reputation. Parameter values are such that, from a project value perspective, the project should be implemented iff $(s_1, s_2) = (s^g, s^g)$, and that the status quo should be maintained in case of conflicting signals (and thus also if $(s_1, s_2) = (s^b, s^b)$).

5.2.1 Theoretical results

From a project-value perspective it is first-best for both members to truthfully reveal their private information within the committee. Undominated voting strategies are then to vote favorably if both messages are positive, and to vote against implementation in the remaining three cases.

If members use the first-best deliberation and voting strategies, then the market infers from implementation that both members have received the same (positive) signal, whereas the decision to maintain the status quo may have resulted from either two concurring negative signals, or from two conflicting ones. As a consequence of the assumed relationship between ability and signals received, conflicting signals are an unambiguous indication that at least one member is of low ability, while observing two concurring signals increases the probability that members are of high ability. As a result, project implementation (which in first best requires two positive signals) yields a higher reputation than maintaining the status quo (which could be due to conflicting signals). If we suppose for the moment that the decision on the project is not accompanied by any statement by the committee, reputations following first-best behavior of the committee members are:

$$\begin{aligned} \widehat{\pi}_{fb}(I=1) &= \frac{1+\pi}{1+\pi^2}\pi > \pi \\ \widehat{\pi}_{fb}(I=0) &= \frac{3-\pi}{3-\pi^2}\pi < \pi. \end{aligned} \tag{5.2}$$

The implication of equations 5.1 and 5.2 is that committee members who care sufficiently about their reputations are willing to give up project payoff for a stronger reputation.[9] Since the loss in project value is smaller in case of conflicting signals than in case of two negative signals, this provides an incentive to committees to implement the project in case of conflicting signals.[10] VS show that there is an upper bound to the likelihood with which the committee wants to implement the project in case of conflicting signals, $\beta = \Pr(I=1|s_1 \neq s_2) < 1/2$. For implementation to be an equilibrium choice in case of

[9] This holds for $\lambda > -p/\left(\frac{1+\pi}{1+\pi^2}\pi - \frac{3-\pi}{3-\pi^2}\pi\right)$.

[10] Note that we write "in case of conflicting signals" rather than "in case of conflicting messages". Indeed, as members' interests are perfectly aligned, truthfully sharing their private information remains part of equilibrium behavior. Hence, messages can be equated with signals.

conflicting signals, the expected loss in project value, $p + E\left[\mu|s^g, s^b\right] = p < 0$, should be compensated by a stronger reputation, $\hat{\pi}\,(I = 1;\beta) > \hat{\pi}\,(I = 0;\beta)$. This requires that $\beta < 1/2$, such that committees with conflicting signals are more likely to maintain the status quo than to implement the project. In other words, in any equilibrium the unconventional decision commands the higher reputation. In equilibrium β is such that:

$$p + \lambda\hat{\pi}\,(I = 1;\beta) = \lambda\hat{\pi}\,(I = 0;\beta)\,. \tag{5.3}$$

Thus, in case of conflicting signals, a member is indifferent between implementing the project and maintaining the status quo. We can rewrite this condition as

$$-p = \lambda\left[\hat{\pi}\,(I = 1;\beta) - \hat{\pi}\,(I = 0;\beta)\right].$$

In this form it is clear that, in equilibrium, changes in λ and/or π leave the difference in reputational utility unaffected and equal to $-p$. The equality also means that in case of two concurring signals, a committee member has a strict preference: project implementation in case of two positive signals, and project rejection in case of two negative signals.

What information is the committee willing to share with the market once a decision has been taken? If members were to truthfully state the messages they exchanged in the meeting or the signals they received, then their reputations would be based on those statements, rather than on the decision taken on the project. Clearly, the larger is the degree of agreement among those messages or signals, the stronger is the reputation of each member. Of course, if the market indeed believes that the committee behaves in this way and adjusts its beliefs accordingly, the best reply of a committee is to present a united front in favor of whatever decision was taken, even if there was no agreement in the meeting. In general, if, for a given decision on the project, the market were to attach a higher reputation to one statement than to another, the committee members will use the former statement, whether it is true or not. The theoretical consequences are clear: the only source of information about the ability of committee members is the decision on the investment project. The investment decision – thanks to its payoff implications – acts as a costly signal. VS draw a second conclusion that is not dictated by game-theoretic logic: members will show a united front. That is, out of the many ways in which a committee can communicate with the outside world that are consistent with equilibrium, they emphasize the equilibrium in which members speak with one voice. This behavior is clearly plausible and finds support in reality.[11] Game theory does dictate that the market should then be able to determine the reputation of a member in the out-of-equilibrium event the committee were not to show a united front. It is consistent with the model to assume that disagreement gives rise to a drop in reputation. The following predictions summarize the above discussion.

Theoretical predictions - Committee behavior

1. Committee members share their private information in the meeting;
2. In case of conflicting signals, the project is implemented with probability β as determined in Equation (5.3);
3. The committee shows a united front to the market in support of the decision it took.

Theoretical predictions - Reputation

1. Implementation commands a higher reputation than maintaining the status quo, irrespective of the value of the parameters;
2. The difference in reputational payoff $\lambda\left[\hat{\pi}\left(I=1;\beta\right)-\hat{\pi}\left(I=0;\beta\right)\right]$ equals $-p>0$, irrespective of λ and π.
3. In equilibrium, the reputation of committee members does not vary with the statements they make. Statements are considered pure cheap talk.

5.3 Experimental design and hypotheses

The experiment closely follows the sequence of events in Figure 5.1. Table 5.1 provides the parameter values used in the experiments, both the ones that are common to all treatments, and the ones that are treatment specific.

At the start of the experiment the computer randomly (and anonymously) matches all subjects to create groups of two. Next, all groups are assigned a role, with half of the groups becoming committees, and half of the groups becoming pairs of bidders,. Each committee is matched with two pairs of bidders and each pair of bidders is matched with two committees. Depending on the number of subjects that shows up for a particular session, committees and pairs of bidders were either connected in the 2×2-scheme as in the left-hand panel of Figure 5.2 or in a $2\times2\times2$-scheme as depicted in the right-hand

[11] Indeed, it would seem odd for a committee to claim that, say, eight out of ten members agreed with the decision taken, if the statement were purely cheap talk. See p. 340 in VS. See also Swank et al. (2008)

Parameter	Value
Common	
Probability high quality investment	0.5
Value of CM	
smart	50
dumb	5
Pay-off of no investment ($X = 0$)	0
Pay-off of investment ($X = 1$)	
high quality	100
low quality	-120
pay-off CM	Investment pay-off + 2nd bids (w_i)
Pay-off bidder	
won auction	Value CM - 2nd bid
lost auction	0
Treatment specific	
Probability CM smart	
high prob. treatment	0.65
low prob. treatment	0.5

Table 5.1: Parameter setting of the experiment

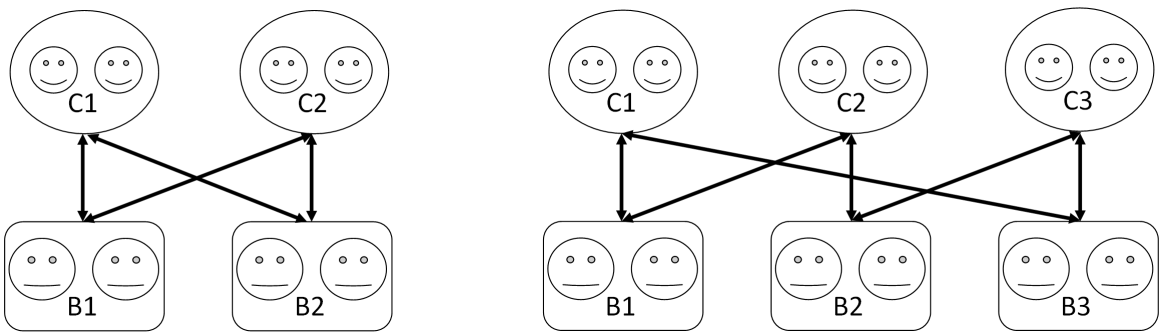

Figure 5.2: Matching schedules

panel of Figure 5.2. Since matching groups do not interact, this allows us to gain up to 2 independent observations in a session with 16 or 20 subjects. The assigned roles and the composition of each committee and pair of bidders remain the same throughout the experiment.

In each round, the sequence of events begins with the computer determining the ability $a_i \in \{\underline{a}, \bar{a}\}$ of each committee member and the state of the world $\mu \in \{-h, h\}$ through a series of random draws. Nature reveals neither a_i nor μ to any subject.

In the low probability treatment, the likelihood that a committee member is highly able equals $\pi_i = 0.5$, while in the high probability treatment it equals $\pi_i = 0.65$. As a subject may consider his ability as something that cannot be manipulated by the experimenter, the instructions read that the likelihood of him getting correct information was 0.5 or 0.65 depending on the treatment, and that his information is random with the complementary likelihood. There is thus no difference with the definition of ability in VS.

After receiving their signals, each committee member can use a chat window for free-form communication with their fellow committee member. This communication is not observed by any other participant in the experiment. Given the information structure of the experiment, theoretically a straw poll would have been sufficient for the required information exchange. However, free-form communication adds realism and has been in experiments on committee decision making before (e.g. Fehrler and Hughes (2014) and Goeree and Yariv (2011)). The free-form communication also allows subjects the possibility to talk about the statements they want to make after investment and thus coordinate on a closed front if they so desire. Examination of the chat logs shows that some committees indeed do.

Next, a committee member votes in favor or against investment. Investment takes place if and only if both members vote favorably. In case of the good state the investment pay-off, $p + \mu$, equals 100 points, while in a bad state $p + \mu$ equals -120 points.

Once both committee members have cast their votes, each committee member observes the investment decision made by his committee and makes a statement to the labor market. This statement indicates his degree of confidence in the decision taken. This captures the essence of the public statements that accompany the committee decisions in VS. Confidence is expressed on a 5-point scale ranging from 'very doubtful' to 'very confident'. Here also a free-form message was possible. However, because we want to show the effect of these statements on the bidding behavior later on, we choose a form of statements that is more easily (and without arbitrary intervention by the experimenters) used in later analysis.

In the (experimental) labor market, our bidders compete to hire individual committee members. The value of a high ability member to a bidder is 50 points, while the value of a low ability member is 5 points. A bidder observes both whether the committee invested in the project and the statements that both members send to the market. However, actual project returns $p+\mu$, signals s_i, messages m_i and votes v_i are not disclosed to the bidders.

The experiment is set up so that each pair of bidders observes two committees and submits four bids, one for the services of each individual committee member in these two committees. As a result, there are four bids for every committee member. To determine which bidder hires which member and at what wage, we use four second-price sealed-bid auctions within each pair of bidders. Thus, in each pair of bidders the bidder submitting the highest bid for a given committee member, wins the services of that member and pays the losing bid. This payment to the committee member is called his wage. Each committee member receives two wages, one from each pair of bidders he is matched with. Bidders were allowed to bid any number between 0 and 50, such that a committee members' wage in any period is between 0 and 100. Ties in the auction are broken at random.

During the experiment, information about the experimental task is presented to the subjects on their screens. At the moment a member receives information about the investment opportunity of his committee, he is reminded of the likelihood of receiving correct information and of the voting rule. At the moment he is prompted to send a confidence statement to the bidders, he is reminded of what other information the bidders would see. Finally, at the end of each round, a member observes the investment payoff, his wage, the sum of his earnings for that round, and the cumulative earnings including that round.

Likewise, at the moment she is prompted to submit bids, each bidder is shown, for both committees she is matched with, the decision taken and the confidence statements of both committee members. She is also reminded of the second-price nature of the bidding procedure and the unconditional probability distributions. At the end of each round, each bidder sees which, if any, auction she has won, and what she and her opponent bid particular committee members and the value of these committee member. In this results screen she also sees her total net pay-off for this round and the cumulative earnings including this round. The position each committee and committee member has on the results screen of a bidder is the same as during the auction. Across rounds, however, these positions are randomly determined, so bidders cannot identify individual committee members or committees over several rounds. This is known to all subjects.

5.3.1 Procedures

The experiments took place in the ESE-econlab at Erasmus University Rotterdam. All subjects were invited via the ESE-econ subject pool using ORSEE, see Greiner (2004). The experiments were run using z-tree, see Fischbacher (2007). Sessions lasted about 1 hour and 45 minutes on average (including instructions and payment). Subjects earned €21 ($28) on average, but there was considerable variation amongst subjects. Each day of experiments two sessions with different treatments were scheduled back to back. The order was determined by the flip of a fair coin.

In all sessions and treatments the structure of the experiment was the same. Upon entering the lab all subjects received verbal instructions about lab protocols and a general description of the experiment. Every subject received a set of written instructions

Variable	Low prob. treatment	High prob. treatment	difference
Gender	.5243902 (.5024781)	.61 (.4902071)	not significant t-test
year	3.402439 (1.961731)	2.66 (1.918649)	0.1% significant χ^2
Risk tolerance	5.536585 (2.121001)	5.52 (1.920017)	not significant t-test
age	21.92683 (3.643972)	21.12 (2.475251)	not significant t-test
econ background	0.9404762 (0.0259705)	0.92 (0.27266)	not significant t-test

Table 5.2: Subject Characteristics

describing the experiment in detail. Subjects were then assigned a seat at random and were given time to read the instructions. The instructions detailed both roles that subjects could play in the experiment, bidders and committee members. The instructions included the particulars of pay-offs and probability distributions and screen-shots of all stages of the experiment, for both roles.[12]

Before the actual experiment began, the computer prompted subjects to answer a set of questions about the pay-offs and probabilities to ensure they understood the set-up. If a subject answered one or more questions incorrectly a pop-up screen would inform him which question he had answered incorrectly, what the correct answer was, and why. The subject was then asked to try again. After all subjects answered all questions correctly, the actual experiment began. In total subjects completed 18 rounds, the first 3 rounds were unincentivized practice rounds. The points earned in the last 15 rounds were converted into Euro's at a fixed exchange rate of €1 per 50 points at the end of the experiment. After these 18 rounds subjects filled out a questionnaire with some background characteristics before getting paid in cash and leaving the lab.

5.4 Experimental Results

5.4.1 Descriptive statistics

In total 184 subjects participated in our two treatments, 84 in the low probability treatment, 100 in the high probability treatment. Half of the subjects were bidders. Each bidder made 4 bids per round for 15 (incentivized) rounds, such that we have 5520 bids, and 2760 wage offers (loosing bids) to work with. Every committee made one investment decision per round while each committee member sent messages, casted votes and made statements individually, yielding 1380 decisions and 2760 messages, votes and statements.

Table 5.2 presents some characteristics of the subjects. Although there is a difference between the two treatments in terms of gender it is not significant. A χ^2 test shows that

[12] A set of instructions can be found in Appendix D.2.

there is significant difference between the two groups in terms of the year of study they are in. The difference stems from a relatively high number of first-year bachelor students in the high prob. treatment and pre-master students in the low prob. treatment. Since all subjects were recruited in the same way, we have no explanation for this difference. The summary statistics of all variables in the database are in the "summary stats" table in the appendix.

All regressions are presented in tables with a similar set-up. The top row lists the dependent variable. Below the horizontal line are the independent variables, with estimated coefficients and the accompanying standard errors in brackets below the coefficients. The bottom of the table shows what, if any, selection is made from the data, as well as the level at which standard errors are clustered, and whether or not period and subject fixed effects are used.

5.4.2 Bidder's behavior - what determines the reputation of a committee member?

The theory of VS leads to three predictions concerning the determinants of reputations, see page 91. In this section we test these predictions using an OLS regression,[13]

$$\begin{aligned} w_{ijt} &= \alpha + \gamma_I D(\text{investment}\,(i,t)) + \gamma_4 D\,(\text{sttmnt}\,(i,t)=4) + \gamma_5 D\,(\text{sttmnt}\,(i,t)=5) \\ &\quad + \gamma_c D\,(\text{closed\_front}\,(i,t)) + FE + \epsilon_{ijt}, \end{aligned}$$

where w_{ijt} is the wage paid by bidder j for committee member i in period t, α a constant, $D(\text{investment}\,(i,t))$ a dummy that is 1 if the committee of which i is a member invested in period t , $D\,(\text{sttmnt}\,(i,t)=4)$ a dummy that equals 1 if member i in period t stated a confidence level of 4 (similarly for $D\,(\text{sttmnt}\,(i,t)=5)$), $D\,(\text{closed\_front}\,(i,t))$ a dummy that equals 1 if the committee i is part of shows a closed front in period t, FE denotes a full set of period and subject fixed effects, and ϵ_{ijt} a zero-mean disturbance term.

The theory predicts, that in both treatments the market pays a higher wage if a committee invests, and second that the difference is related to the expected loss in project value in case of conflicting signals, $-p = 10$. Note that the experiment is set-up in such a way that committee members receive wages from two bidders, one from each pair of bidders that is matched with a committee. Thus the difference in wages between

[13] Since the bidders play a second-price auction the highest bid should theoretically contain an upward bias, while the second price should be an unbiased estimate of the value of the item. We therefore only report the regressions based on the second price, the wage, in each auction. Running the regression on all bids does not qualitatively effect the outcomes of the regressions.

investment and no investment should equal 5, $\gamma_I = 5$, in both treatments. Finally, the theory predicts that in either treatment the wage is not affected by statements that committee members make, $\gamma_4 = \gamma_5 = \gamma_c = 0$.

Table 5.3 presents convincing evidence for the first two predictions.

VARIABLES	(1) Wage	(2) Wage	(3) Wage	(4) Wage	(5) Wage	(6) Wage
investment	6.642***			4.855**	6.097***	5.568***
	(1.179)			(1.980)	(1.176)	(1.058)
Sttmnt=4		1.047		-1.101	2.114	0.421
		(1.257)		(1.264)	(1.709)	(1.123)
Sttmnt=5		2.303**		0.344	-0.0285	0.176
		(0.969)		(1.055)	(1.334)	(0.923)
closed_front			6.849***	5.198***	5.954***	
			(1.141)	(1.540)	(1.030)	
high_sttmnt						-3.163***
						(0.907)
low_sttmnt						-7.960***
						(1.173)
Constant	16.72***	18.31***	16.93***	14.36***	15.11***	20.39***
	(1.178)	(1.581)	(1.341)	(1.666)	(2.351)	(1.413)
Observations	2,760	2,760	2,760	1,260	1,500	2,760
R-squared	0.397	0.340	0.389	0.383	0.399	0.442
Subject FE	Yes	Yes	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
cluster	labgroup	labgroup	labgroup	labgroup	labgroup	labgroup
Treatment	H&L	H&L	H&L	L	H	H&L

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5.3: Formation of reputation

From the table it is clear that investment increases the reputation of committee members. F-tests do not reject the hypothesis that the coefficient equals 5 in each treatment. Thus, as predicted by the theory, the difference in wage is such that the members of a committee that receives conflicting signals are (made) indifferent between investing and not investing: on average, what a member gains in terms of wage thanks to investing equals what he loses in terms of expected project payoff.

The table also shows that statements made by committee members influence the wages paid: a united front and higher statements lead to higher wages. This goes against the equilibrium prediction of the theory. In sections 5.4.3 and 5.4.4 we investigate possible

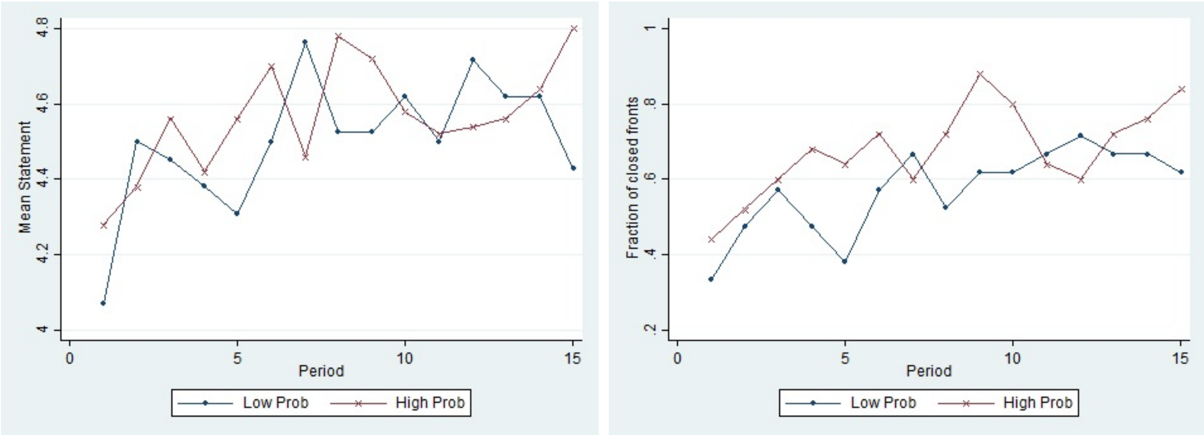

Figure 5.3: Evolution of Statements and Closed fronts

reasons for this finding. Repeating regressions (4) and (5) with interaction terms between *closed_front* and *investment*, and between *closed_front* and the statement variables did not change much. All interaction terms where negative and neither the size nor the significance of the other coefficient was influenced much. The increase in wage thanks to these statement-related variables seems to be independent of the investment decision. Also note that the effects of the closed front and high statements are somewhat hard to disentangle as statements become both higher and more likely to be aligned over time, see Figure 5.3. Since 5 is the highest possible statement, high statements are more likely to concur. Despite this issue, table 5.3 clearly indicates that making statements lower than 5 would cause the bids to decrease. With an average statement above 4.5, it is clear that committee members realize this. Similarly a committee that presented a closed front received wages about 6 points higher than committees that did not. To ensure that this effect is not caused by the ceiling effect, we create two dummies. The dummy *high_sttmnt* (*low_sttmnt*) is set to 1 if the statement made by a committee member was strictly higher (strictly lower) than the statement of his fellow committee member in the same period. Regression 5 shows that if a committee does not show a closed front, both the member making the relatively high statement and the one making the relatively low statement receive a penalty for the apparent disagreement. The behavior of the market should cause committee members to concur on the highest statement. Looking at Figure 5.3 it seems that committee members indeed picked up on this incentive.

Signals / Statements	-1	0	1
-1	518	104	6
1	5	98	501

Table 5.4: Statements as a function of signals

	Investment	No Investment
high prob	43%	57%
low prob	34%	66%

Table 5.5: Investment prob. per treatment

5.4.3 Committee behavior

We already observed in the previous section that, in case of conflicting signals, what a member gains in terms of wage thanks to investing equals what he loses in terms of expected project payoff. This is an important aspect of optimal committee behavior.

The theory generates three more predictions concerning the behavior of a committee. First, it predicts that members share their private information in the meeting. To see whether this is indeed the case, we coded all the messages exchanged in the chat box. Messages explicitly talking about the signal received are coded 1 if the committee member claims to have received a positive signal, and -1 if he claims a negative signal. Messages that only state "invest" or "we should invest" are coded as 0, since it is unsure whether it is a suggested course of action, or whether it pertains to the signal received.[14] For the purpose of this analysis every conversation between two committee members in a particular round was coded as two messages, one for each committee member.

Table 5.4 shows that, by and large, members indeed shared their private information. Only 11 out of the 1232 coded messages contained an outright lie (all in different committees), while 1019 messages directly reveal the true signal received. Both the lies and the messages we coded as 0's seem to be distributed equally over positive and negative signals. This suggests that they are not used strategically within the committee.

A second prediction deals with the relationship between signals and the likelihood of investment. In either treatment, a committee should invest in case of two positive signals and should not invest in case of no positive signals. In case of conflicting signals, investment should happen with a positive probability, and this probability should be higher in the high probability treatment. On the whole, the a priori unlikely decision - investment - should remain the decision that is taken less often.

[14] Appendix D.1 shows the conversion table.

# positive signals received	0	1	2
High prob. treatment	9.8	25	97.6
Low prob. treatment	1.162	24.1	77.1

Table 5.6: Investment prob. as a function of signals received

To start with the latter, Table 5.5 reports the incidence of either decision in both treatments. In both treatments *investment* was indeed the least likely decision, while a binomial probability test rejects the hypothesis that the investment happens at least 50% of the time (pr<0.000 for low prob, pr<0.01 high prob).

Table 5.6 shows the percentage of committees that invest as a function of the number of positive signals its members received.

Clearly, the *additional* investment in the high probability treatment stems from committees with zero or two positive signals, not from committees with mixed signals as would be predicted by theory. This is confirmed by a logit regression (unreported) on the effect of the treatment on the likelihood of investing for the sample of committees that received conflicting signals. The coefficient on the dummy for a high-probability treatment is positive, but small and statistically insignificant. Note, however, that in theory and in the data a committee with conflicting signals is indifferent between investing and not investing. For the given difference in wage that bidders pay, any probability of investment is a best-reply. What is not consistent with optimizing behavior is the 23% of committees in the low probability treatment that received two positive signals but decided not to invest.

A final prediction of the theory is that committees show a united front to the outside world. We have seen in figure 5.3 that there is indeed a strong tendency to show united fronts and express high levels of confidence in the decision taken. Furthermore, both the confidence and the percentage of closed fronts increases with practice. Committee members seem to learn over time to report the same, highly confident message, especially in the high probability treatment. Nevertheless, it is also clear from the bidders' behavior that some committees leave money on the table by not showing a united front. Although the committees' strategy is not a perfect best response to bidder behavior, in general they seem to pick up on the incentives generated by the "reputational market". In the next subsection we take a closer look at how these responses change over time to see if committees learn.

VARIABLES	(1) Investment	(2) Closed front
sum_sign=1	4.719***	-1.027***
	(0.350)	(0.200)
sum_sign=2	10.18***	0.346
	(0.464)	(0.237)
Period	0.0935***	0.196***
	(0.0271)	(0.0169)
sum_sign=1 X Period	-0.212***	-0.0748***
	(0.0318)	(0.0220)
sum_sign=2 X Period	-0.160***	-0.110***
	(0.0408)	(0.0232)
Investment		1.142***
		(0.125)
Observations	5,280	5,280
Number of labsubjects	88	88
cluster	NONE	NONE
Treatment	H&L	H&L
Period FE	NO	NO
Subject FE	YES	YES

Standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5.7: Committees learning

Do committees learn?

The theory predicts that committees perfectly take into account the market responses to their behavior. However, neither markets nor bidders are perfectly playing equilibrium behavior. One explanation, based on figure 5.3, could be that subjects learn to play equilibrium over time. We focus on *Investment* and *Closed_front* since we know from table 5.3 that the reputation of committee members increases most in these variables, and hence learning should be most pronounced. In table 5.7 we show the effect of time by examining how the decision making process of committee members changes over time. Investment becomes marginally more likely over time for committees that have two negative signals, while it becomes mildly less likely for committees with 1 or 2 positive signals. Overall, however, these coefficients are dwarfed by the coefficients on *sum_sign=1* and *sum_sign=2*. Even in the 15th, and last, period the effect of the signals is as expected: most investment occurs if both signals are positive etc.

	high prob	low prob
smart	71.11%	66.45%
dumb	61.45%	48.13%

Table 5.8: Percentage of Closed Ranks

Committees do appear to learn to present a closed front. The coefficient on *Period* is larger than either interaction term, indicating that all committees were showing more closed fronts over time. Interestingly enough, the learning effect seems strongest in committees that received 0 positive signals, even though table 5.10 indicates that it is the group with mixed signals that is least likely overall to present a united front, and thus has "most to learn". The positive coefficient on *Investment* is somewhat puzzling, it appears that investing committee members are less likely to admit internal conflict. Although game theory provides no reason why this would be so, perhaps having money on the line through investment makes it more important for them to appear convinced or united (recall that most closed fronts occur for the highest stated confidence).

5.4.4 Do bidders use the available information efficiently?

In section 5.4.2 we noted that bidders pay higher wages if they observe a united front and if committee members express high levels of confidence in the decision taken. This finding is in conflict with the equilibrium prediction that there is no relevant information about ability in the statements made. Are bidders irrational or is there relevant information about ability in these statements? If the latter is the case, do bidders make the most of this information, i.e., is their behavior a best reply? The analysis in this section shows that statements indeed contain useful information and that bidders make use of that information.

Tables 5.8, 5.9 and 5.10 show the marginal distributions of the cheap talk statements as a function of ability and the signals received by the committee. Thus, the bottom right cell of table 5.8 indicates that in the low probability treatment 48.3% of low ability committee members make the same statement as their fellow committee member.

In both treatments it appears that high ability members are more likely to make the same statements. Markets should therefore increase their bids for committees that show a closed front.

If we look at the distribution of statements over the 5 possible levels of confidence, it is quite clear that high ability members more often indicate that they are "very confident". The difference in distributions is significant in a χ^2-test ($p < 0.000$).

Confidence level	1	2	3	4	5	total
High ability	1.38	1.63	4.76	14.04	78.2	100
Low ability	3.44	3.09	10.82	18.56	64.09	100

Table 5.9: Statements made

Both observations could be caused by the signal distributions. A high ability member is more likely to have the same signal as his fellow committee member. His committee is therefore more likely to take a decision in accordance with his signal. To show the strength of this explanation table 5.10 displays the number of closed fronts as a function of the number of positive signals received in a committee. A χ^2-test rejects the independence

# of positive signals	0	1	2	total
different Sttmnt	136	254	122	512
Closed front	320	228	320	868
total	456	482	442	1380

Table 5.10: Signals and Closed fronts

of these distributions (p=0.00): committee members are more likely to present a closed front if they have concurring signals.

These findings imply that statements contain relevant information about ability. Committee members that presented a closed front and/or made high statements were more likely to be of high ability. The question then becomes whether bidders made efficient use of this information.

Theoretically, the wage, the price paid in the second-price auction, should equal the estimated value of the services of a committee member, i.e., $w_{ijt} = E(v_{it}|X_{ijt})$, where v_{it} is the value of member i in period t, and X_{ijt} contains all information bidder j has on committee member i in period t. Using the properties of OLS predictors as conditional expected values, we could run two separate regressions, one for w_{ijt}, $w_{ijt} = \alpha_1 + \gamma X_{ijt} + \eta_{ijt}$, and one for v_{it}, $v_{it} = \alpha_0 + \beta X_{ijt} + \epsilon_{it}$. Bidders use information efficiently if $\beta = \gamma$. In that case, the wage changes in exactly the same way as the expected value of committee members as a function of observed covariates. Any systematic deviations of this equality would indicate inefficient use of information by the bidders. However, the fact that ability influences not only v_{it} but also the observables in X_{ijt} gives rise to a missing variable and/or endogeneity problem. Using these two regressions might cause undesirable biases. We therefore did two things. To check for missing variable bias we include the actual value v_{it} of a committee member as a regressor in table 5.11. The insignificance of the coefficient on v_{it} indicates that the X_{ijt} capture all information used by the bidders.

To deal with the (theoretical) correlation between v_{it} and w_{ijt} we created the variable *Difference*, $\Delta_{ijt} = v_{it} - w_{ijt}$. This variable measures any difference between the value of a committee member and his reputation, such that it should capture any systematic deviations from the efficient use of information. In both treatments *Difference* runs from -45 till 50, which are the minimum and maximum values implied by the constraint on the bids and the value distribution. The variable has an average value of around 5 and a standard deviation of roughly 25 points. Figure 5.4 shows a histogram of the empirical distribution. The distribution clearly shows the peaks caused by the bi-modal value distribution, implying we should be careful with the standard-errors in the regression. Because the standard errors are clustered in all regressions below, this should not cause large over-estimations of the precision of the coefficients.

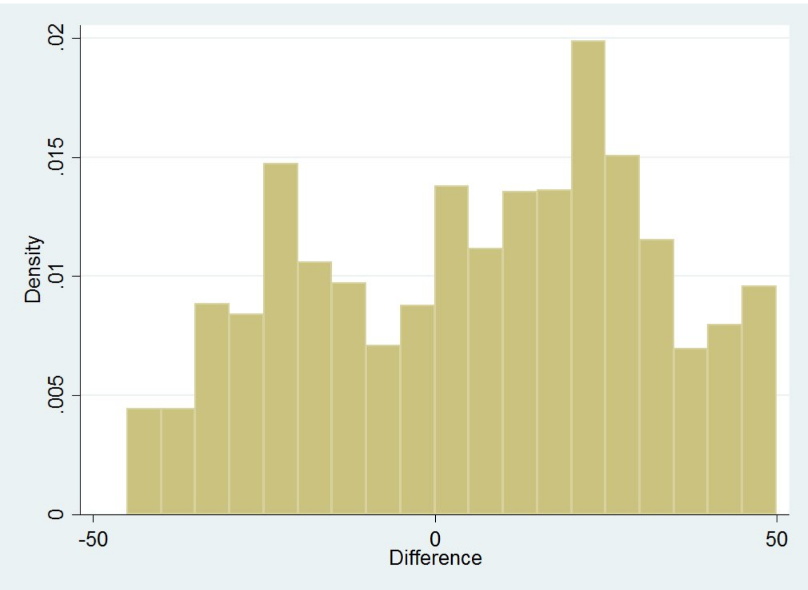

Figure 5.4: Distribution of Difference

We can now run the suggested test, in the form of a single regression orthogonality test:

$$\begin{aligned}\Delta_{ijt} &= \alpha + \gamma_I I(i,t) + \gamma_4 D(\text{sttmnt}(i,t)=4) + \gamma_5 D(\text{sttmnt}(i,t)=5) \\ &\quad + \gamma_c D(\text{closed\_front}(i,t)) + FE + \epsilon_{ijt},\end{aligned}$$

with ϵ_{ijt} a zero mean error term. The strategy of a bidder is a best-reply, and bidders use the available information efficiently if $\gamma_I = \gamma_4 = \gamma_5 = \gamma_c = 0$.

Furthermore, to ensure we capture the individual decisions we include both period and bidder fixed effects in the regressions. In the final regressions we try to explain the effect of the bidder FE on bids, by comparing the regressions with and without FE and background variables. Table 5.11 shows the outcomes by treatment.

Before discussing the outcomes, first note that the standard errors in the regressions in table 5.11 are clustered on the matching group level instead of the individual level. If bidders use information efficiently, the theoretical predictions are identical to the statistical null-hypothesis: all coefficients are zero. Clustering too finely would increase the standard errors too much, which would make the orthogonality test too likely to accept the null hypothesis. Since bidders and committee members could influence each other's behavior, not clustering would be falsely assuming full independence between the observations. The largest unit on which committees and bidders could influence each other was the matching group. This level of clustering is the most critical test available that still corrects for correlation in the actions of individual subjects, and therefore we cluster the standard errors at this level.

Regressions 1 and 2 in Table 5.11 show that individual bidders make quite efficient use of all observable information: all estimated coefficients are insignificant. The small size of the coefficients indicate that, even without clustering standard errors, there is no reason to think that investment is over- or undervalued as a sign of ability by the bidders. The coefficients on the message and closed-front variables are larger in absolute size, such that the clustering potentially has more influence on the statistical significance. In fact, if standard errors are not clustered, *sttmnt=4* and *closed_front* become significant at the 10% level in the high probability treatment. It seems that some extra information in the message variables could have gone unnoticed by the bidders in that treatment. This pattern is also visible in an F-test for joint significance of the regressors. In the low probability treatment this yields a $p = 0.71$, clearly indicating that these variables are not likely to carry much importance. In the high probability treatment the same test yields a $p = 0.086$, indicating that there is some indication that bidders missed some information in that treatment. Overall, however, these effects seem quite marginal, bidders appear to have been quite effective in weighing the information presented to them.

The findings in regression 3 and 4 seem to strengthen the argument that information is used efficiently. The coefficient on *Value* is insignificant in both regression 3 and 4. The effect of the unobservable random variable seems to be captured entirely by the visible aspects of committee behavior. Thus we feel quite certain in concluding that bidders' strategies are (close to) best-replies and the variables *Investment*, *closed_front*, *Sttmnt=4* and *Sttmnt=5* contain the information used by the bidders to determine their bids.

VARIABLES	(1) Difference	(2) Difference	(3) Wage	(4) Wage
Investment	-0.597	-1.587	4.960*	6.118***
	(5.159)	(1.772)	(2.657)	(1.133)
closed_front	1.644	-2.716	5.366**	5.969***
	(3.527)	(2.386)	(2.071)	(1.189)
Sttmnt=4	3.104	-4.974	-1.052	2.100
	(3.186)	(2.812)	(1.132)	(1.772)
Sttmnt=5	2.386	1.229	0.411	-0.0228
	(1.931)	(2.401)	(1.132)	(1.291)
Value			-0.0245	-0.00472
			(0.0173)	(0.0162)
Constant	2.029	10.89**	14.76***	15.24***
	(4.202)	(4.827)	(1.908)	(2.377)
Observations	1,260	1,500	1,260	1,500
R-squared	0.108	0.092	0.384	0.399
cluster	Match group	Match group	Match group	Match group
Treatment	L	H	L	H
Period FE	Yes	Yes	Yes	Yes
Subject FE	Yes	Yes	Yes	Yes

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$

Table 5.11: Efficient use of information

5.5 Robustness analysis

5.5.1 Bidders comparing committees

In the theory of VS, agents are assumed to know the game and the equilibrium strategies completely and thoroughly. In our experiment, after subjects had read the instructions, their knowledge was tested through a questionnaire. To ensure subjects are aware of the relevant information, the computer program continuously presented relevant distributions and aspects of the game to the subject during the experiment. Besides this information, the bidders submit their bids through a screen that allows them to compare the two committees they are bidding on. From the game-theoretic point of view, comparing behavior of two committees whose investment opportunities are unrelated makes no sense. Behaviorally, however, there is the possibility that bidders compare the behavior of each committee with that of the other committee, as if they are comparing the behavior with equilibrium strategies. This would, for instance, imply that the increase in wage following investment of any committee depends negatively on the investment of the other committee. If bidders follow this behavioral pattern, a bid should be higher (i) for the committee that invests if the other does not invest; (ii) for the committee that shows a united front if the other does not; (iii) if a member makes a statement that is higher than the average statement made by the other committee.

VARIABLES	(1) Wage	(2) Wage
Sttmnt=4	-0.730	2.335
	(1.085)	(1.730)
Sttmnt=5	1.156	0.446
	(0.932)	(1.238)
investment	5.901**	7.840***
	(2.131)	(1.376)
other_not_invest	3.542***	1.471
	(0.914)	(1.041)
Invest. X other_not	-2.634	-3.034
	(2.574)	(1.868)
closed_front	3.129*	5.853***
	(1.560)	(1.245)
other_not_closed	0.597	1.811
	(1.103)	(1.130)
closed X other_not	3.172*	-0.409
	(1.546)	(1.968)
extra_confident	2.704	2.117
	(2.343)	(2.750)
Constant	11.85***	13.14***
	(1.881)	(2.412)
Observations	1,260	1,500
R-squared	0.411	0.406
Subject FE	Yes	Yes
Period FE	Yes	Yes
Cluster	individual	individual
Treatment	L	H

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5.12: Contrasting two committees

To test these hypotheses we created several variables that keep track of the information on the bidders' screen. The variable *other_not_invest* is a dummy that is set to one if the other committee on the bidders screen does not invest, similarly *other_not_closed* equals 1 if the other committee does not show a closed front. The variable *extra_confident*, measures the difference between the statement made by a committee member and the average of the statements made by the members of the other committee. If the bidders compare the committees on their screen these variables, or the interaction between these variables and the behavior of the committee members, should have a significant effect on the wage levels.

Overall the results in table 5.12 seem to give some credence to the idea that bidders compare committees, mostly in the low probability treatment. If we compare the regressions of the low prob. treatment with the regressions in table 5.3, it seems that part of the increase in wages following a closed front or investment is caused by the comparison with the other committee on the bidders' screen. In the high probability treatment no such effects are found. In both treatments and both tables *investment* and *closed_front* have a positive effect, such that the comparison with other committees does not actually change the incentive effects created by career concerns. Interestingly the effect of statements made disappears completely in table 5.12 compared to table 5.3. As predicted by theory the effect of cheap talk statements on wages is not very robust. Although the information on a bidders' screen about another committee can have some effect on the wage of a particular committee member, the main predictions from VS still carry through. An expert that cares about his reputation should invest more than first-best and always make sure the members of his committee present a closed front.

5.5.2 Background variables

To get a better idea of the which committee members were in favor of investment we present a logit regression that relates the voting behavior of each subject with his or her background variables and some situation variables. The results are shown in table 5.13.

As expected the amount of positive signals increased votes in favor of investment. Committee members are more likely to vote in favor of investment in the high probability treatment, consistent with the larger amount of investment in that treatment. The effect of *age* on voting behavior is economically insignificant. The coefficient is quite small and our subjects do not differ much in age, the youngest being 18 and the oldest 30 while most subjects are between 19 and 24. Our measure of willingness to take risks, *risk_prox*, does not affect voting behavior of committee members in general, it does help predict whether or not a committee member votes in favor of investment in case the committee has mixed signals. Subjects that are more willing to take risk, are also more likely to invest in a uncertain project to increase their reputation. Since this is a risky investment, this effect is as expected. A similar regression on bidding behavior showed no significant

	(1)	(2)
VARIABLES	vote	Bid
_Itreatment_2	0.596**	5.088***
	(0.302)	(1.798)
risk_prox	-0.00743	-0.262
	(0.0945)	(0.466)
gender	-0.0788	2.520
	(0.355)	(1.718)
econ_bg	0.434	2.332
	(0.537)	(2.647)
year	0.0473	-0.165
	(0.0821)	(0.580)
age	0.0657**	-0.222
	(0.0323)	(0.369)
_Isum_sign_1	1.836***	
	(0.265)	
_Isum_sign_2	5.301***	
	(0.452)	
_IsumXrisk__1	0.207**	
	(0.101)	
_IsumXrisk__2	0.196	
	(0.179)	
_IQMsg2_2		0.331
		(1.609)
_IQMsg2_3		1.012
		(1.430)
decision		5.075***
		(0.834)
closed_front		4.662***
		(1.001)
Constant	-4.160***	23.76***
	(0.892)	(7.390)
Observations	5,520	2,728
R-squared		0.195
Subject FE	NO	NO
Period FE	Yes	Yes
cluster	individual	individual

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5.13: Effect of background variables

effects of subject background variables on the bids made. The insignificant effect of year is quite reassuring given the difference in the distribution between the two treatments. In general the background variables seem to have little impact on the qualitative results of our experiment.

5.6 Conclusion

In this paper, we have reported on an experiment in which committee members cared both about their reputations and about the returns on the project. The control in the lab allows us to discern when project value is maximized and when career concerns cause subjects to distort their choices. This level of observability allows us to test the theoretical mechanisms predicted by the career concerns literature. Our results indicate that the model of VS captures the reputational mechanisms accurately enough to predict their consequences. This allows us to predict theoretically when and how career concerns can have undesirable consequences, as well as allow us to use theory to design better decision making procedures.

5.6.1 Future research

In both our treatments our expert committees have career concerns, it would be useful to compare our findings with those of a situation in which committee members do not care about their reputations. We can then study how in the absence of career concerns committee members deal with the uncertainty concerning the profitability of investment and how markets evaluate the ability of committee members. The theory of VS predicts that the difference in perceived ability is higher in the absence of career-driven behavior and that a committee invests only if both its members have positive private information.

It is common in experimental research to allow inexperienced subjects to "learn to play the game." Besides providing instructions, quizzes and incentives, the most important element that should favor learning is repetition: letting subjects act repeatedly within the same situation. An alternative, or complementary, way of stimulating learning relies not on repetition but on decomposition. Rather than letting subjects experience the full-fledged version of the game, subjects can be allowed to play first a simple version, to which additional layers of complexity are being added over time. In this paper, we followed the standard practice. An experimental analysis of the theory of VS also allows for the second approach. Subjects can first gain experience with a situation in which career concerns play no role. This not only simplifies the objective function of committee members. It also gets rid of the interaction between committee members and markets. Thus, the task of markets – to evaluate committee members – should also become easier, as their actions does not influence the behavior of the committee members. In a second step, committee members can be made to care about both the investment and their careers. In fact, this incremental approach is also used in VS to explain the theory as well as in the theory section above. It would be interesting to see whether the two ways of stimulating learning yield different outcomes, and whether one type is more conducive than the other in bringing subjects to equilibrium play.

In our experiment committee members are forced to make a statement about their confidence, while this might not always be the case in the field. Although it seems that subjects use the statements strategically, a stronger test of the strategic nature could, for instance, be to allow members to refrain from making a statement and/or make the statement itself, but not its content, costly. If experts still make a statement, it seems more likely that they are using it strategically.

Appendix A

Appendices When a Price is Enough

For bookkeeping, the Jacobian of first derivatives $\phi'(\cdot)$ of any function $\phi(\cdot) : \mathcal{R}^a \to \mathcal{R}^b$, is of dimension $b \times a$, while the Hessian of second derivatives $\phi''(\cdot)$ is of dimension $ab \times a$. For any multi-vector functions $\psi(\mathbf{z}^1, \mathbf{z}^2, \ldots) : \mathcal{R}^{a^1} \times \mathcal{R}^{a^2} \ldots \to \mathcal{R}$ the vector of first derivatives ψ_{z^i} is of dimension $a^i \times 1$ and the matrix of second derivatives $\psi_{z^i z^j}$ is of dimension $a^i \times a^j$ where the dimension of the matrix follows the order of the subscripts. In addition, let superscript T be the transpose operator. Vectors and multi-dimensional constructs are denoted in bold, scalars are in normal font.

A.1 Proof of Lemma 2.3

Proof. Due to non-satiation of the utility function we know that the budget constraint will hold with equality such that we know that:

$$y^*(\mathbf{n}) = q(\mathbf{x}^*(\mathbf{n})) - T(\mathbf{x}^*(\mathbf{n}))$$

Direct substitution of the budget constraint into the utility function allows us to write the first-order conditions to problem (2.2) as:

$$\mathbf{0} = u_{\mathbf{x}} + (q' - T')^T u_y \tag{A.1}$$

which directly implies equations (2.7) and (2.8).

Now take the second-order derivative of the utility function with respect to $\mathbf{x}$ to get the second-order conditions:

$$u_{\mathbf{xx}}+\left(2u_{\mathbf{x}y}+u_{yy}\left(q'\left(\mathbf{x}^{*}\right)-T'\left(\mathbf{x}^{*}\right)\right)^{T}\right)\left(q'\left(\mathbf{x}^{*}\right)-T'\left(\mathbf{x}^{*}\right)\right)+u_{y}\left(q''\left(\mathbf{x}^{*}\right)-T''\left(\mathbf{x}^{*}\right)\right)\lessdot 0 \tag{A.2}$$

Differentiate the marginal rate of substitution, $\mathbf{s}$, toward $\mathbf{x}$, using the definition of $\mathbf{s}$ and the implicit function theorem to define $y\left(u,\mathbf{x},\mathbf{n}\right)$, to get:

$$\frac{\partial\mathbf{s}\left(\mathbf{x},y\left(u,\mathbf{x},\mathbf{n}\right),\mathbf{n}\right)}{\partial\mathbf{x}}=-\frac{\left(u_{\mathbf{xx}}+2u_{\mathbf{x}y}\mathbf{s}^{T}\right)-u_{yy}\mathbf{s}\mathbf{s}^{T}}{u_{y}} \tag{A.3}$$

Now combining (2.8) with (A.3) allows us to simplify (A.2) and obtain the third condition:

$$\begin{aligned}
-\left(\frac{\partial\mathbf{s}\left(\mathbf{x},y\left(u,\mathbf{x},\mathbf{n}\right),\mathbf{n}\right)}{\partial\mathbf{x}}+q''\left(\mathbf{x}^{*}\right)-T''\left(\mathbf{x}^{*}\right)\right)u_{y} &\lessdot 0\Leftrightarrow\\
-\frac{\partial\mathbf{s}\left(\mathbf{x},y\left(u,\mathbf{x},\mathbf{n}\right),\mathbf{n}\right)}{\partial\mathbf{x}}|_{\{\mathbf{x}=\mathbf{x}^{*}(n),y=y^{*}(n)\}}+q''\left(\mathbf{x}^{*}\right)-T''\left(\mathbf{x}^{*}\right) &\lessdot 0
\end{aligned}$$

where the final step follows from the assumption that $u_y > 0$.

The last condition of lemma 2.3, follows directly from the fact that the maximum has to be the best option in the entire choice-space, while conditions i.) to iii.) only guarantee it attains a local maximum in $\{\mathbf{X},Y\}^{*}$.

□

A.2 Proof to Proposition 2.4 and corollary 2.5

Proof. Start by assuming that condition iv.) of lemma 2.3 holds. Suppose taxes are equated to wedges for a, presumed optimal, incentive-compatible allocation. Consider a joint deviation of an agent from the bundle assigned to him when the original allocation is incentive compatible, resource feasible, and within this set on the Pareto frontier. Denote the joint deviation by $\alpha\Delta\mathbf{x}$ where $\alpha > 0$ and $\Delta\mathbf{x}$ is a $k\times 1$ vector which we normalize to length one, on the market by a particular agent of type $\mathbf{m}$. The budget constraint (2.5) defines the amount of y the deviating agent receives as:

$$y=q\left(\mathbf{x}^{*}\left(\mathbf{m}\right)+\alpha\Delta\mathbf{x}\right)-T\left(\mathbf{x}^{*}\left(\mathbf{m}\right)+\alpha\Delta\mathbf{x}\right).$$

Note that the resulting allocation after the joint deviation is incentive compatible since the original allocation was incentive compatible and the exchange occurs at market prices. Through the utility function, the utility gain, $\mathcal{U}\left(\alpha\boldsymbol{\Delta}\mathbf{x}\right)$, of such a deviation can therefore be written as:

$$\mathcal{U}\left(\alpha\mathbf{\Delta x}\right)=u\left(\mathbf{x}^{*}\left(\mathbf{m}\right)+\alpha\mathbf{\Delta x},q\left(\mathbf{x}^{*}\left(\mathbf{m}\right)+\alpha\Delta\mathbf{x}\right)-T\left(\mathbf{x}^{*}\left(\mathbf{m}\right)+\alpha\Delta\mathbf{x}\right),\mathbf{n}\right)-u^{*}\left(\mathbf{m}\right),$$

where u^{*} denotes the utility the agent obtains without deviation. By a second-order Taylor approximation around $u^{*}\left(\mathbf{m}\right)$ the utility gain can be written as:

$$\begin{aligned}
\mathcal{U}\left(\alpha\mathbf{\Delta x}\right) &\approx \left(u_{\mathbf{x}}^{T}+u_{y}\left(q'-T'\right)\right)\alpha\mathbf{\Delta x}+\\
&\quad \frac{1}{2}\alpha^{2}\mathbf{\Delta x}^{T}\left(u_{\mathbf{xx}}+\left(2u_{\mathbf{x}y}+u_{yy}\left(q'-T'\right)^{T}\right)\left(q'-T'\right)+u_{y}\left(q''-T''\right)\right)\mathbf{\Delta x},\\
&\approx \frac{1}{2}\alpha^{2}\mathbf{\Delta x}^{T}\left(u_{\mathbf{xx}}+\left(2u_{\mathbf{x}y}+u_{yy}\left(q'-T'\right)^{T}\right)\left(q'-T'\right)+u_{y}\left(q''-T''\right)\right)\mathbf{\Delta x},\\
&\approx \frac{1}{2}\alpha^{2}u_{y}\mathbf{\Delta x}^{T}\left(-\frac{\partial\mathbf{s}\left(\mathbf{x},y\left(u,\mathbf{x},\mathbf{m}\right),\mathbf{m}\right)}{\partial\mathbf{x}}|_{\{\mathbf{x}=\mathbf{x}^{*},y=y^{*},u=u^{*}\}}+q''\left(\mathbf{x}^{*}\right)-T''\left(\mathbf{x}^{*}\right)\right)\mathbf{\Delta x},
\end{aligned}$$

where the first order terms equals zero by the assumption that taxes are equated to wedges, $T_{i}'\left(\mathbf{x}^{*}\left(\mathbf{m}\right)\right)=\mathcal{W}_{i}(\mathbf{m})$, and we have used equation (A.3) for simplification.

Now suppose second-order implementation constraints (2.9) are not satisfied such that:

$$\mathcal{H}=-\frac{\partial\mathbf{s}\left(\mathbf{x},y\left(u,\mathbf{x},\mathbf{m}\right),\mathbf{m}\right)}{\partial\mathbf{x}}|_{\{\mathbf{x}=\mathbf{x}^{*},y=y^{*},u^{*}\}}+q''\left(\mathbf{x}^{*}\right)-T''\left(\mathbf{x}^{*}\right),$$

is not negative semi-definite. In that case there exists at least one deviation strategy $\alpha\Delta\hat{\mathbf{x}}$ which yields a positive utility gain: $\mathcal{U}\left(\alpha\mathbf{\Delta}\hat{\mathbf{x}}\right)>0$. To see this note that if this were not the case, by definition $\mathcal{H}$ would be negative semi-definite. In addition, since $\mathcal{H}$ is a matrix of second-order derivatives and therefore symmetric, the opposite deviation $-\alpha\mathbf{\Delta}\hat{\mathbf{x}}$ yields approximately the same utility gain. A second-order Taylor approximation then yields: $\mathcal{U}\left(\alpha\mathbf{\Delta}\hat{\mathbf{x}}\right)\approx\mathcal{U}\left(-\alpha\mathbf{\Delta}\hat{\mathbf{x}}\right)$. Since, the approximation error of the Taylor expansion limits to zero as $\alpha\rightarrow0$, there must exist at least two deviation strategies $\alpha\mathbf{\Delta}\hat{\mathbf{x}}$, and $-\alpha\mathbf{\Delta}\hat{\mathbf{x}}$ that yield positive utility for sufficiently small α.

The change in tax revenue due to such a deviation can also be approximated by means of a second-order Taylor expansion:

$$\begin{aligned}
\mathcal{T}\left(\alpha\Delta\hat{\mathbf{x}}\right) &= T\left(\mathbf{x}^{*}\left(\mathbf{n}\right)+\alpha\Delta\hat{\mathbf{x}}\right)-T\left(\mathbf{x}^{*}\left(\mathbf{n}\right)\right)\\
&\approx \alpha T'\Delta\hat{\mathbf{x}}+\frac{1}{2}\alpha^{2}\Delta\hat{\mathbf{x}}^{T}T''\Delta\hat{\mathbf{x}}.
\end{aligned}$$

The first-order term will always be non-negative for deviation strategy $\alpha\Delta\hat{\mathbf{x}}$ or for -$\alpha\Delta\hat{\mathbf{x}}$. If for one of the choices, the first-order term is positive, the first-order term dominates the second-order term for sufficiently small α. Hence, there is a deviation that results in higher tax revenue. If on the other hand the first-order term is zero, we need to consider the second-order term. If it is negative the tax schedule contains an internal maximum in at least one vector subspace of $\mathbf{X}$ which violates the assumption stated in the proposition. Therefore, if the first-order term is zero the second term must be non-negative. Hence, tax revenue always weakly increases in either $-\alpha\Delta\hat{\mathbf{x}}$ or $\alpha\Delta\hat{\mathbf{x}}$ for sufficiently small α. It follows that at least one of the two deviations is resource feasible.

Therefore, under the conditions stated in our proposition, if the tax system is not successful in implementing the allocation on the Pareto frontier, there exist at least one deviation which increase utility of at least one agent, leaves the allocation resource feasible and incentive compatible. The utility of all other agents remains unaffected, since there are no externalities by assumption, and is therefore a Pareto improvement. Hence, we run into a contradiction, since a Pareto improvement cannot exist over an allocation that is already on the Pareto frontier proving the main part of the proposition.

The stated corollary follows from the fact that the second-best allocation of a welfarist planner always lies on the Pareto frontier. To see this note that if the allocation were not on the Pareto frontier, by definition the planner could increase the utility of at least one agent without decreasing the utility of any other agent. Such a deviation increases the value of the planner's objective function, and hence the original allocation could not have been the second best (this is also shown in Werning, 2007,Brendon, 2013).

Such that if the tax schedule off the allocation is set in such a way the maximum is a global one and condition iv.) of lemma 2.3 holds, the tax implements the allocation.

□

A.3 Proof to Proposition 2.9

Proof. Note that in the direct mechanism, reporting your type to the planner is equivalent to choosing a bundle $\{\mathbf{x}, y\}$ among the set of bundles that are assigned to a type, $\{\mathbf{X}, Y\}^*$. Hence, we can rewrite the optimization problem for the agents in the direct mechanism (see definition 2.1) as follows:

$$\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\} = \arg\max_{\mathbf{x},y} \{u(\mathbf{x}, y, \mathbf{n}) : \{\mathbf{x}, y\} \in \{\mathbf{X}, Y\}^*\}.$$

In addition, by equation (2.7), for each bundle of $\mathbf{x}^*(\mathbf{n})$ the tax system provides the corresponding level $y^*(\mathbf{n})$. Hence, if the tax system, $T(\mathbf{x})$, satisfies the equation (2.7) we can further simplify the optimization problem to:

$$\begin{aligned}\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\} &= \arg\max_{\mathbf{x},y} \{u(\mathbf{x}, y, \mathbf{n}) : y = q(\mathbf{x}) - T(\mathbf{x}), \mathbf{x} \in \mathbf{X}^*\}, \\ &= \arg\max_{\mathbf{x},y} \{u(\mathbf{x}, y, \mathbf{n}) : y = q(\mathbf{x}) - T(\mathbf{x}), \mathbf{x} \in \mathbf{X}\},\end{aligned}$$

where the final step follows from the fact that the allocation is surjective such that the set of assigned bundles is equal to the set of all available bundles on the market: $\mathbf{X}^* = \mathbf{X}$. This is exactly equal to optimization problem 2.2, and hence, the tax system that satisfies (2.7) can implement any surjective allocation. Note that we do not need to check whether the tax system also satisfies equations (2.8), (2.9) or condition iv.). Since we have assumed that the original allocation is incentive compatible, it follows immediately that these constraints are also satisfied provided the allocation is surjective.

Finally, note that the tax system that implements a surjective allocation is unique. Equation (2.7) implicitly defines the value of the tax function for all $\mathbf{x} \in \mathbf{X}^*$:

$$T(\mathbf{x}^*(\mathbf{n})) = q(\mathbf{x}^*(\mathbf{n})) - y^*(\mathbf{n}).$$

Therefore, equation (2.7) uniquely determines the value of $T(\mathbf{x})$ for all $\mathbf{x} \in \mathbf{X}^*$. Since by surjectiveness $\mathbf{X}^* = \mathbf{X}$ it defines the value of $T(\mathbf{x})$ over its entire domain. Hence, there is only one function $T(\mathbf{x})$ that satisfies (2.7). Since (2.7) is a necessary condition for implementation, it follows that there is only one tax system that can implement a surjective allocation. □

A.4 Example

A.4.1 The Second-Best Allocation

Our graphical example in section 4 is based on the following model. Suppose the economy is inhabited by couples and their children. Each couple j derives utility from their joint consumption, which is assumed to be the untaxed numeraire commodity Y^j. In addition, the couples derive disutility from hours worked by each spouse i in couple j, l_i^j, as well as disutility from the interaction between hours worked, $l_1^j l_2^j$. This gives couples an incentive for partial specialization where one partner works more hours, while the other spends more time at home. For simplicity we assume preferences for labor are symmetric, such that the couple does not have a preference for which of the two spouses supplies the highest amount of labor. Each partner i in couple j is assumed to have identical ability, n^j. This could be the result of perfect assortative matching. As is standard in the literature, we assume gross income is the product of hours worked and ability: $X_i^j = n^j l_i^j$. We assume the following functional form for the utility function of the parents:

$$\begin{aligned} U_p\left(Y, l_1, l_2, l_1 l_2\right) &= \gamma_Y \log\left(Y\right) - \frac{\gamma_l}{\alpha} \sum_{i=1}^{2} l_i^{\alpha} - \frac{1}{\beta}\left(l_1 l_2\right)^{\beta}, \\ &= \gamma_Y \log\left(Y\right) - \frac{\gamma_X}{\alpha} \sum_{i=1}^{2} \left(\frac{X_i}{n}\right)^{\alpha} - \frac{1}{\beta}\left(\frac{X_1 X_2}{n^2}\right)^{\beta}, \end{aligned}$$

where utility is assumed to be logarithmic in consumption of Y, CES with parameter $\alpha > 1$ in each spouse's hours worked, and CES with parameter β in the interaction of hours worked. γ_Y and γ_X are scaling parameters which allows one to weight the relative importance of each of the three terms in the utility function.

Each couple is assumed to have one child. Their utility function does not weight the utility of the child.[1] The child does not make any decisions but its utility is affected by the decisions of his parents. This creates a classic externality where the preferences of the child are not included in the transactions made by the parents, but do affect total welfare. The child is assumed to receive utility of consumption of the household. In addition, it receives disutility from ±labor, since leisure time of each parent is spend with the child. Finally, the child enjoys spending time with both of his parents, and therefore receives utility if the parents do not specialize but share the task of rearing him. We capture this by letting the child's utility increase in $l_1 l_2$. We assume the following utility function for the child's utility function:

$$U_c\left(Y, l_1, l_2, l_1 l_2\right) = \delta_Y \log\left(Y\right) - \frac{\delta_l}{\xi} \sum_{i=1}^{2} l_i^{\xi} + \frac{1}{\zeta}\left(l_1 l_2\right)^{\zeta},$$

where $\xi > 1$ is the parameter for CES disutility of labor, ζ the parameter for utility from the interaction of hours worked, and δ_Y and δ_l scaling parameters.

In the simulations we take a discrete approximation of our model, such that n follows a discrete distribution with J possible outcomes $n_{\min}, \ldots, n_{\max}$.[2] We assume each group has similar size f such that the mass of a particular type is equal to $g = \frac{1}{J}$. We assume the welfare function of the planner is a concave increasing function of joint family utility

[1] Note that the main intuition carries over to the case where the parents take account of some, but not all of the utility of the child.

[2] Note that our model formally does not apply to discrete models of taxation. However, as we increase the number of types, we can approximate the case fior continuous type with increasing precision.

$U = U_P + U_C$. Hence, social welfare can be written as:

$$SW = \frac{1}{f}\sum_{j=1}^{J}\frac{1}{1-\nu}U_j^{1-\nu}, \tag{A.4}$$

where ν measures the rate of relative inequality aversion.

We normalize all prizes to one such that the resource constraint of the economy can be written as:

$$\begin{aligned} 0 &\geq \frac{1}{f}\sum_{j=1}^{J}Y^j - X_1 - X_2, \\ 0 &\geq \sum_{j=1}^{J}Y^j - X_1^j - X_2^j. \end{aligned} \tag{A.5}$$

In the direct mechanism, the planner is assumed to offer bundles $\left(X_1^j, X_2^j, Y^j\right)$ to each household j on the basis of a their reported type. The allocation is incentive compatible if each couple truthfully reports their type to the planner:

$$j = \arg\max_{k=1,\ldots,J} U_p\left(Y^k, \frac{X_1^k}{n^j}, \frac{X_2^k}{n^j}, \frac{X_1^k X_2^k}{\left(n^j\right)^2}\right) \quad \forall \quad j = 1, \ldots, J. \tag{A.6}$$

Note that incentive compatibility only depends on the utility of the parents, and not of the children since they do not take any decision. In the simulations we take the first-order approach and calculate the optimum under the condition that each couple j has at least as high utility by telling the truth as by mimicking type $j-1$:

$$U_p\left(Y^j, \frac{X_1^j}{n^j}, \frac{X_2^j}{n^j}, \frac{X_1^j X_2^j}{\left(n^j\right)^2}\right) \geq U_p\left(Y^{j-1}, \frac{X_1^{j-1}}{n^j}, \frac{X_2^{j-1}}{n^j}, \frac{X_1^{j-1} X_2^{j-1}}{\left(n^j\right)^2}\right) \quad \forall \quad j = 2, \ldots, J. \tag{A.7}$$

At the end of the simulation we check whether the allocation calculated under the first-order approach also satisfies (A.6) and in the reported simulation this is indeed the case. The optimal allocation was calculated numerically by maximizing equation (A.4) subject to (A.7) and (A.5).

A.4.2 Market Implementation

The market implementation is the standard Mirrlees implementation as described in the end of section 3. That is, the budget constraint of the households is given by:

$$Y = X_1 - X_2 - T(X_1) - T(X_2),$$

and marginal tax rates are found by equating taxes to the optimal wedges in the second-best allocation, while assuming separability of the marginal tax rates.

A.4.3 Parametrization

For the figures in the main text we used the following parameterization: $\alpha = \beta = \xi = \zeta = 1.5$. In addition, we chose $\gamma_Y = \gamma_l = \delta_Y = \delta_l = 1$. Finally, the planner is assumed to be utilitarian such that $\nu = 0$. This yields the following set of equations:

$$\begin{aligned}
SW &= \frac{1}{f}\sum_{j=1}^{J}\frac{1}{1-\nu}U_j^{1-\nu} \\
U &= U_p + U_c \\
U_p &= \log(Y) - \frac{1}{1.5}\sum_{i=1}^{2}\left(\frac{X_i}{n}\right)^{1.5} - \frac{1}{1.5}\left(\frac{X_1X_2}{n^2}\right)^{1.5}, \\
U_c &= \log(Y) - \frac{1}{1.5}\sum_{i=1}^{2}\left(\frac{X_i}{n}\right)^{1.5} + \frac{1}{1.5}\left(\frac{X_1X_2}{n^2}\right)^{1.5}, \\
0 &\geq \sum_{j=1}^{J} Y^j - X_1^j - X_2^j.
\end{aligned}$$

We programmed this optimization problem in Matlab. Codes are available upon

Appendix B

Appendices As Easy as ABC?

B.1 Proof of Proposition 3.1

Proof. The first order condition for incentive compatibility, in a local maximum, is given by:

$$\begin{aligned}\mathbf{0}_p &= \frac{\partial u\left(\mathbf{x}^*\left(\mathbf{m}\right), y^*\left(\mathbf{m}\right), \mathbf{n}\right)}{\partial \mathbf{m}}|_{\mathbf{m}=\mathbf{n}}, \\ &= \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T u_{\mathbf{x}}\left(\mathbf{x}^*\left(\mathbf{n}\right), y^*\left(\mathbf{n}\right), \mathbf{n}\right) + u_y\left(\mathbf{x}^*\left(\mathbf{n}\right), y^*\left(\mathbf{n}\right), \mathbf{n}\right) y^{*\prime}\left(\mathbf{n}\right)^T, \end{aligned} \tag{B.1}$$

where $\mathbf{0}_p$ denotes a p-column vector of zeros. This can be rewritten to:

$$y^{*\prime}\left(\mathbf{n}\right) = s\left(\mathbf{x}^*\left(\mathbf{n}\right), y^*\left(\mathbf{n}\right), \mathbf{n}\right)^T \mathbf{x}^{*\prime}\left(\mathbf{n}\right),$$

proving equation (3.3).

We can derive (3.5) from (3.2) using the envelope theorem:

$$\begin{aligned}V'\left(\mathbf{n}\right) &= \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T u_{\mathbf{x}} + u_y y^*\left(\mathbf{n}\right)^T + u_{\mathbf{n}}^T, \\ V'\left(\mathbf{n}\right) &= u_{\mathbf{n}}\left(\mathbf{x}^*\left(\mathbf{n}\right), y^*\left(\mathbf{n}\right), \mathbf{n}\right)^T, \end{aligned} \tag{B.2}$$

where the latter equality follows from the first-order conditions.

The second-order conditions of a local maximum are :

$$\frac{\partial^2 u\left(\mathbf{x}^*\left(\mathbf{m}\right), y^*\left(\mathbf{m}\right), \mathbf{n}\right)}{\partial \mathbf{m}^2}|_{\mathbf{m}=\mathbf{n}} \lessdot 0, \tag{B.3}$$

Where $\preccurlyeq 0$ denotes the negative semi-definiteness of the matrix.
Taking the derivative of (B.2) with respect to $\mathbf{m}$ gives:

$$\begin{aligned}\frac{\partial^2 u\left(\mathbf{x}^*\left(\mathbf{m}\right),y^*\left(\mathbf{m}\right),\mathbf{n}\right)}{\partial \mathbf{m}^2} &= \left(u_{\mathbf{x}}\left(\mathbf{x}^*\left(\mathbf{m}\right),y^*\left(\mathbf{m}\right),\mathbf{n}\right)^T \otimes I_p\right)\mathbf{x}^{*''}\left(\mathbf{m}\right) \\ &+\mathbf{x}^{*\prime}\left(\mathbf{m}\right)^T u_{\mathbf{xx}}\left(\mathbf{x}^*\left(\mathbf{m}\right),y^*\left(\mathbf{m}\right),\mathbf{n}\right)\mathbf{x}^{*\prime}\left(\mathbf{m}\right) \\ &+u_{yy}\left(\mathbf{x}^*\left(\mathbf{m}\right),y^*\left(\mathbf{m}\right),\mathbf{n}\right)y^{*\prime}\left(\mathbf{m}\right)y^{*\prime}\left(\mathbf{m}\right)^T \\ &+u_y\left(\mathbf{x}^*\left(\mathbf{m}\right),y^*\left(\mathbf{m}\right),\mathbf{n}\right)y^{*''}\left(\mathbf{m}\right). \end{aligned} \tag{B.4}$$

where $\otimes$ denotes the Kronecker product.

To simplify this expression we take the total derivative of the first order condition (B.1):

$$D_{\mathbf{n}}\mathbf{0}_p = D_{\mathbf{n}}\left[\mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T u_{\mathbf{x}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right) + u_y\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)y^{*\prime}\left(\mathbf{n}\right)^T\right] \tag{B.5}$$

$$\begin{aligned}\mathbf{0}_{p\times p} &= \left(u_{\mathbf{x}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)^T \otimes I_p\right)\mathbf{x}^{*''}\left(\mathbf{n}\right)^T + \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T u_{\mathbf{xx}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)\mathbf{x}^{*\prime}\left(\mathbf{n}\right) + \\ &u_{yy}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)y^{*\prime}\left(\mathbf{n}\right)y^{*\prime}\left(\mathbf{n}\right)^T + u_y\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)y^{*''}\left(\mathbf{n}\right) + \\ &\mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T u_{\mathbf{xn}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right) + y^{*\prime}\left(\mathbf{n}\right)^T u_{y\mathbf{n}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right), \end{aligned} \tag{B.6}$$

Now combine equations (B.5,B.3,B.4) to get the following expression:

$$\begin{aligned} 0 &\succcurlyeq \left(u_{\mathbf{x}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)^T \otimes I_p\right)\mathbf{x}^{*''}\left(\mathbf{n}\right)^T + \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T u_{\mathbf{xx}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)\mathbf{x}^{*\prime}\left(\mathbf{n}\right) + \\ &u_{yy}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)y^{*\prime}\left(\mathbf{n}\right)y^{*\prime}\left(\mathbf{n}\right)^T + u_y\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right)y^{*''}\left(\mathbf{n}\right) - \mathbf{0}_{p\times p} \\ 0 &\succcurlyeq \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T u_{\mathbf{xn}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right) + y^{*\prime}\left(\mathbf{n}\right)^T u_{y\mathbf{n}}\left(\mathbf{x}^*\left(\mathbf{n}\right),y^*\left(\mathbf{n}\right),\mathbf{n}\right). \end{aligned} \tag{B.7}$$

Then partially differentiate the vector of shadow prices with respect to $\mathbf{n}$ to get:

$$\begin{aligned} s_{\mathbf{n}} &= \frac{-u_{\mathbf{xn}}u_y + u_{\mathbf{x}}u_{y\mathbf{n}}}{\left(u_y\right)^2} \\ u_{\mathbf{xn}} &= -s_{\mathbf{n}}u_y - su_{y\mathbf{n}}, \end{aligned}$$

and substitute this result and (3.3) into (B.7) to yield

$$\begin{aligned} 0 &\succcurlyeq \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T\left(-s_{\mathbf{n}}u_y - su_{y\mathbf{n}}\right) + y^{*\prime}\left(\mathbf{n}\right)^T u_{y\mathbf{n}} \\ 0 &\preccurlyeq \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T s_{\mathbf{n}}u_y \Leftrightarrow \\ 0 &\preccurlyeq \mathbf{x}^{*\prime}\left(\mathbf{n}\right)^T s_{\mathbf{n}}, \end{aligned}$$

where the final inequality, equation (3.4), follows from the fact that $u_y > 0$.

An equivalent expression can be derived by totally differentiating, equation (B.2) with respect to $\mathbf{n}$:

$$\begin{aligned} V''(\mathbf{n}) &= Du_{\mathbf{n}}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right) \\ &= \mathbf{x}^{*\prime}(\mathbf{n})^T u_{\mathbf{xn}}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right) + y^{*\prime}(\mathbf{n})^T u_{y\mathbf{n}}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right) \\ &\quad + u_{\mathbf{nn}}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right) \end{aligned}$$

Now combine this last expression with (B.7) to get the final equation:

$$V''(\mathbf{n}) - u_{\mathbf{nn}}\left(\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n}), \mathbf{n}\right) \gg 0.$$

□

B.2 Proof of Proposition 3.3

Proof. Starting from the first order conditions (3.15), (3.14) and (3.16). First solving (3.14) for yields η:

$$\eta = \frac{\lambda f + u_{y\mathbf{n}}\theta}{u_y}.$$

Now substitute this expression into (3.15) and simplify to get the desired equation:

$$\begin{aligned} \lambda q'^T f - u_{\mathbf{xn}}\theta + \frac{\lambda f + u_{y\mathbf{n}}\theta}{u_y} u_{\mathbf{x}} &= \mathbf{0}_k, \\ \lambda q_{x_i} f - u_{x_i\mathbf{n}}\theta + \frac{\lambda f + u_{y\mathbf{n}}\theta}{u_y} u_{x_i} &= \mathbf{0}_k, \\ \lambda f\left(q_{x_i} - s_i\right) &= u_{x_i\mathbf{n}}\theta + u_{y\mathbf{n}}\theta\mathbf{s}, \\ q_{x_i} - s_i &= \left(u_{x_i\mathbf{n}} + u_{y\mathbf{n}} s_i\right)\frac{\theta}{\lambda f}, \\ q_{x_i} - s_i &= \sum_{j=1}^{p} -\frac{\partial s_i}{\partial n_j}\theta_j \frac{u_y}{\lambda f}, \\ \frac{q_{x_i} - s_i}{s_i} &= \sum_{j=1}^{p} \varepsilon_{x_i n_j} \times \theta_j \frac{u_y}{\lambda} \times \frac{1}{n_j f}. \end{aligned}$$

□

B.3 Proof of Proposition 3.7

Proof. Proposition 3.7 consists of 2 statements. First, we prove that the separating partition $\mathbf{N}_S$ is either convex or empty.

Define the types α, β, γ s.t. $\alpha, \gamma \in \mathbf{N}_S$ and $\exists k \in (0,1)$ s.t. $(1-k)\,\alpha + k\gamma = \beta$. Denote by $\{\widetilde{x}, \widetilde{y}\}$ the solution to the full problem, and by $u\left(\{\overline{x}, \overline{y}\}, \mathbf{n}\right) = \max_{\{x,y\}} u(x, y, \mathbf{n} | \{\widetilde{\mathbf{x}}, \widetilde{\mathbf{y}}\}^{\leftarrow} \subseteq \mathbf{N}_B)$ the bunching choice that delivers type $\mathbf{n}$ the highest utility from the set of bunching allocations.

Then define the function $L(\mathbf{n}) = u\left(\{\mathbf{x}^*(\mathbf{n}), y^*(\mathbf{n})\}, \mathbf{n}\right) - u\left(\{\overline{x}, \overline{y}\}, \mathbf{n}\right)$. By individual rationality we know that $0 < L(\alpha)$,$L(\gamma)$. Equation (3.6) implies that L is convex and monotone, and continuity of u implies L is piece-wise continuous. For any $0 < k < 1$ it must be the case that $0 < L(\ \beta)$, and β must also be part of the separating set.

Second we prove that bunching occurs below the boundary $\partial N_S \cap\ \partial N_B$, and only one such boundary exists.

Since $V(\mathbf{n})$ is continuous, individual rationality requires any type at $\mathbf{b} \in \partial N_S \cap\ \partial N_B$, $u\left(\{\mathbf{x}^*(\mathbf{b}), y^*(\mathbf{b})\}, \mathbf{b}\right) = u\left(\{\overline{x}, \overline{y}\}, \mathbf{b}\right)$. While on the separating set equation (3.5) guarantees that the first derivative of u has to be equal to the first-order derivative of v. By equation (3.6), however, the second-order derivative on the optimal allocation has to be higher than the second derivative of the utility function on $\{\overline{x}, \overline{y}\}$, such that the utility profile of the separating and the bunching partition cross only once for each type, and they cross at $\{\mathbf{x}^*(\mathbf{b}), y^*(\mathbf{b})\}$ for type $\mathbf{b}$. The single-crossing property (3.4) then implies that this is the unique border and that separation occurs on side of the border where types have relatively high utility.

This can easily be seen if we assume $\frac{\partial u}{\partial n_i} > 0 \quad \forall i$. In that case, for any type $\mathbf{g}$, with $b_j \leq g_j\ j \in \{1, .., p\}$ and at least one inequality strict, the utility profile of the optimal allocation, $V(\mathbf{n})$, has to be higher than $u\left(\{\overline{x}, \overline{y}\}, \mathbf{g}\right)$. For any type $\mathbf{g}$, therefore, (3.5) holds and the allocation found above the boundary $\partial N_S \cap\ \partial N_B$ induces separation. Simultaneously, below the boundary equation (3.6) cannot hold, which (together with (3.5) and continuity of V) implies that bunching yields a higher utility, such that bunching occurs there. □

Appendix C

Appendices Who dares?

C.1 The Questionaire

1. What is your current professional role?

O Supervisory board member
O CEO, president or director
O CFO or controller
O COO
O Board member
O Consultant
O Manager
O Analyst

2 In which sector does the company operate in which you are active?

O Automotive
O Construction / Materials
O Chemistry
O Retail / Wholesale
O Services
O Electronics
O Engineering
O Pharmacy
O Financial institutions
O Information
O Media
O Metal
O Non-profit
O Oil / Mining
O Education / Science
O (Semi-)government
O Paper / Packaging
O Telecommunications
O Transport
O Utilities (gas / electric / other)
O Food and beverages

3 What are your company's annual revenues?

O Less than € 50 million
O Between € 50 and € 100 million
O Between € 100 and € 500 million
O Between € 500 million and € 1 billion
O More than € 1 billion

4 How many employees does your company have?

- O Less than 50
- O Between 50 and 100
- O Between 100 and 1,000
- O Between 1,000 and 10,000
- O More than 10,000

Your background

5 How often were you involved in making the following investment decisions in the last fifteen years?

	Never	1 to 5 times	5 to 10 times	10 to 15 times	More than 15 times
Expansion into new markets	○	○	○	○	○
Expansion of production capacity	○	○	○	○	○
An innovation or R&D process	○	○	○	○	○
An IT project	○	○	○	○	○
An acquisition	○	○	○	○	○

6 What is the typical size (in % of annual turnover of the company) of the following investment decisions that you have dealt with?

	Not applicable	1%	2%	3%	4%	5%-10%	10%-20%	20%-30%	30%-40%	40%-50%	More than 50%
Expansion into new markets	○	○	○	○	○	○	○	○	○	○	○
Expansion of production capacity	○	○	○	○	○	○	○	○	○	○	○
An innovation or R&D process	○	○	○	○	○	○	○	○	○	○	○
An IT-project	○	○	○	○	○	○	○	○	○	○	○
An acquisition	○	○	○	○	○	○	○	○	○	○	○

Scenario 1

7 Suppose you are asked to assess the following investment opportunity:

In a strategy meeting it became clear that expansion of your business is possible by establishing a new marketing channel abroad. Your marketing and sales department estimates the cost at € [amount] million. If this plan fails, your company loses this amount, but there is no further damage. Whether the channel can fulfill the high expectations is still unclear. Your marketing and sales department estimates the probability of success to be [success rate].

What should be the minimum net present value (the value of the revenues after deduction of the investment discounted back to today) of this investment, would you, in your role, agree with this investment?

8 How risky do you consider this investment?

	Riskless			Neither riskless nor risky			Very risky
I find this investment	○	○	○	○	○	○	○

Scenario 2

9 Suppose you are asked to assess the following investment opportunity:

Your company considers taking a big step in its development by up scaling the production. This involves a total investment of € [amount] million. If it turns out that the up-scaling is not maintainable in the future, your business will lose this investment, but there is no further damage. The probability that up-scaling is successful is currently estimated to be [success rate].

What should be the minimum net present value (the value of the revenues after deduction of the investment discounted back to today) of this expansion of capacity, would you, in your role, agree with this investment?

10 How risky do you consider this investment?

	Riskless			Neither riskless nor risky			Very risky
I find this investment	○	○	○	○	○	○	○

Scenario 3

11 Suppose you are asked to assess the following investment opportunity:

Your R&D department suggests establishing a new manufacturing technique within the current conduct of business. To find out whether this technique is applicable, research has to be done. If it turns out that the new technique is not applicable, your company loses the research costs, but there is no further damage. Development costs are estimated at € [amount] million. The probability that the production technology will be applicable it is estimated to be [success rate].

What should be the minimum net present value (the value of the revenues after deduction of the investment discounted back to today) of this project, would you, in your role agree with this investment?

12 How risky do you consider this investment?

	Riskless			Neither riskless nor risky			Very risky
I find this investment	○	○	○	○	○	○	○

Scenario 4

13 Suppose you are asked to assess the following investment opportunity:

The IT department of your company considers using a different system, which in time can lead to a significant savings in costs. The costs for the introduction of the system are € [amount] million. If it appears that the new system cannot fulfill the expectations, your company loses this amount, but there is no further damage. Whether the system can fulfill the high expectations, is still unclear. Your IT department currently estimates this probability to be [success rate].

What should be the minimum net present value (the value of the revenues after deduction of the investment discounted back to today) of this IT-project, would you, in your role, agree with this investment?

14 How risky do you consider this investment?

	Riskless			Neither riskless nor risky			Very risky
I find this investment	○	○	○	○	○	○	○

15 How do you rate yourself: Are you in general a person who is willing to take risks or are you someone who tries to avoid risks?

Please fill in the scale where a value of 0 means "Not at all willing to take risks" and the value 10 means "very willing to take risks".

0: Not at all willing to take risks	1	2	3	4	5	6	7	8	9	10: Very willing to take risks
○	○	○	○	○	○	○	○	○	○	○

16 Are you, in your role in the company, willing to take risks or do you try to avoid risks?

0	1	2	3	4	5	6	7	8	9	10
○	○	○	○	○	○	○	○	○	○	○

17 How do you assess an average CEO, president or director in his or her role: Is that a person who is willing to take risks or is that someone who tries to avoid risks?

0	1	2	3	4	5	6	7	8	9	10
○	○	○	○	○	○	○	○	○	○	○

18 How do you assess an average CFO or controller in his or her role: Is that a person who is willing to take risks or is that someone who tries to avoid risks?

0	1	2	3	4	5	6	7	8	9	10
○	○	○	○	○	○	○	○	○	○	○

19 How do you assess an average non-executive (like a commissioner or regulator) in his or her role: Is that a person who is willing to take risks or is that someone who tries to avoid risks?

0	1	2	3	4	5	6	7	8	9	10
○	○	○	○	○	○	○	○	○	○	○

20 What is your main background?
- O Entrepreneur
- O Specialist (among which accountants, doctors and lawyers)
- O Finance
- O Marketing/Sales
- O Operations
- O Other

21 Did you use a calculator to answer this questionnaire?
- O Yes
- O No

22 What is your date of birth?

23 What is your gender?
- O Male
- O Female

C.2 Summary Statistics

Table C.1: Summary Statistics

1. What is your current professional role?

O Supervisory board member	15.1 %
O CEO, president or director	16.8 %
O CFO or controller	20.9 %
O COO	1.0 %
O Board member	2.7 %
O Consultant	15.1 %
O Manager	22.6 %
O Analyst	5.1 %

2. What is your company's annual revenue?

O Less than €50 million	28.4 %
O Between €50 and €100 million	11.0 %
O Between €100 and €500 million	18.8 %
O Between €500 million and €1 billion	8.2 %
O More than €1 billion	32.9 %

3. How many employees does your company have?

O Less than 50	14.7 %
O Between 50 and 100	8.9 %
O Between 100 and 1,000	29.5 %
O Between 1,000 and 10,000	25%
O More than 10,000	10.5 %

Note: Due to missing values, the percentages do not add up to 100%

	mean	**std.dev**	**min**	**max**
Used calculator	0.12	0.3262	0	1
Private willingness to take risks	7.07	1.75	2	11
age	46	11	24	68
confidence	0.847	1.46	-2	5

Appendix D

Appendices Career Concerned Managers

D.1 Message coding

Coding messages	
1	Good
	Good here
	It's good
	My info is good
	My info says invest
	Investment good here
	Good advice
	Mine is good
	good news
	I have good
	Ik heb goed
	Goed
-1	Bad here
	Bad
	It's bad
	My info is bad
	My info says not invest
	My one says bad
	Investment bad
	Bad news
	I have bad
	Ik heb slecht
	Slecht
0	Good?
	Bad?
	Invest
	Not invest
	Yes
	Yes?
	I prefer not invest
	I prefer invest
	I would go for not invest
	I would go for invest
coding depends on message sent within group before that	Me too
	Same here

Table D.1: Coding of messages in chat box

D.2 Experiment Instructions

Instructions

The experiment will consist out of 18 rounds, 3 trial rounds and 15 paid rounds, followed by a questionnaire. The first three rounds are trial rounds and do not influence your payment, so you can use them to get to know the game and the environment. From the 4th round (1st non-trial round) on your decisions will influence your payment. The top left box of the screen always displays what round you are in and whether or not it is a trial round.

At the start of the experiment you will be randomly linked to another player. Each of these pairs is then randomly assigned a role. Half of the pairs will be committee members; the other half will be bidders. Both the player you are matched with and the role you are assigned at the beginning of the experiment remain the same during the experiment.

During the experiment you can earn points. At the end of the experiment all earned points will be converted into Euros at a fixed exchange rate and will be paid to you by the experimenter. The exchange rate will be shown on your computer screen during the experiments and will remain constant during the experiment.

Below the tasks of both types of players are explained with screen shots. In the screen shots some areas are marked for explanatory purposes. The marks do not appear in the actual experiments. The values of different variables in the screenshots are examples and you can (and will) see different values during the experiment.

Note that during the game you will only see the screens belonging to your role, and not the other screens. Before the actual experiment starts you will get a small quiz to check your understanding of the game. After the quiz your role will be randomly drawn by the computer and the game begins.

For ease of notation the instruction only uses male pro-nouns.

The tasks: committee members

Each round, each committee gets the opportunity to invest. The investment opportunity is either good or bad in quality both with 50% probability. The investment decision influences the committee members' pay-offs. At the end of the round the investment pay-offs are calculated by the computer: not investing yields 0 with certainty, while investing yields 100 points if it is a good investment and -120 points if it is a bad investment. The quality of an investment opportunity does not depend on the quality of any other investment opportunity within the same round or within another round. The investment opportunity is presented on a screen like the one below (without the colored numbers):

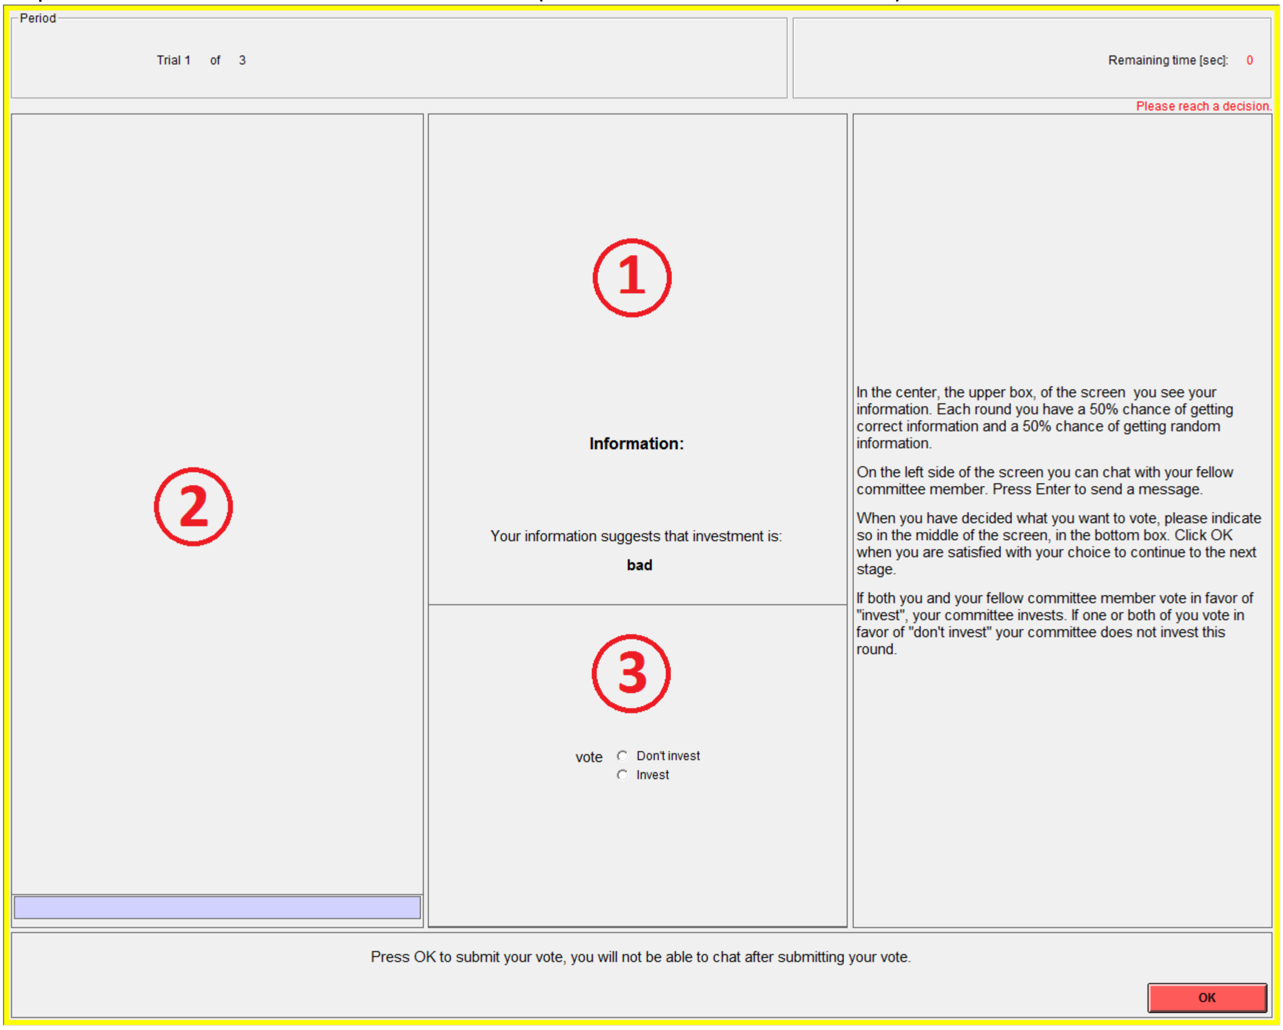

In the center of the screen (indicated with 1) each committee member sees his own information. Each round each individual committee member has a 50% chance of getting correct information and a 50% chance of getting random information. Correct information suggests that the investment opportunity is "good" when it is actually good and it suggests that the investment is "bad" when the investment opportunity is bad. Random information suggests either "good" or "bad" at random with a 50% chance for either, independent of the quality of the investment. A committee member does not know whether he has received correct information or random information. Every round all information is renewed, so having correct information in one round does not influence the chance of having correct information in the next round.

During this stage the committee members can communicate through the chat window on the left side of the screen (indicated with 2). Communication is private, so only the two committee members in each committee can see what they are saying to each other.

The committee members have to vote in favor of "Invest" or "Don't invest" (indicated with 3). If both committee members vote in favor of "invest", the investment takes place. If one or both committee member(s) vote(s) in favor of “Don’t invest”, the investment does not take place this round. Each round a new investment opportunity is created for all committees.

After voting, all committee members go to the next stage:

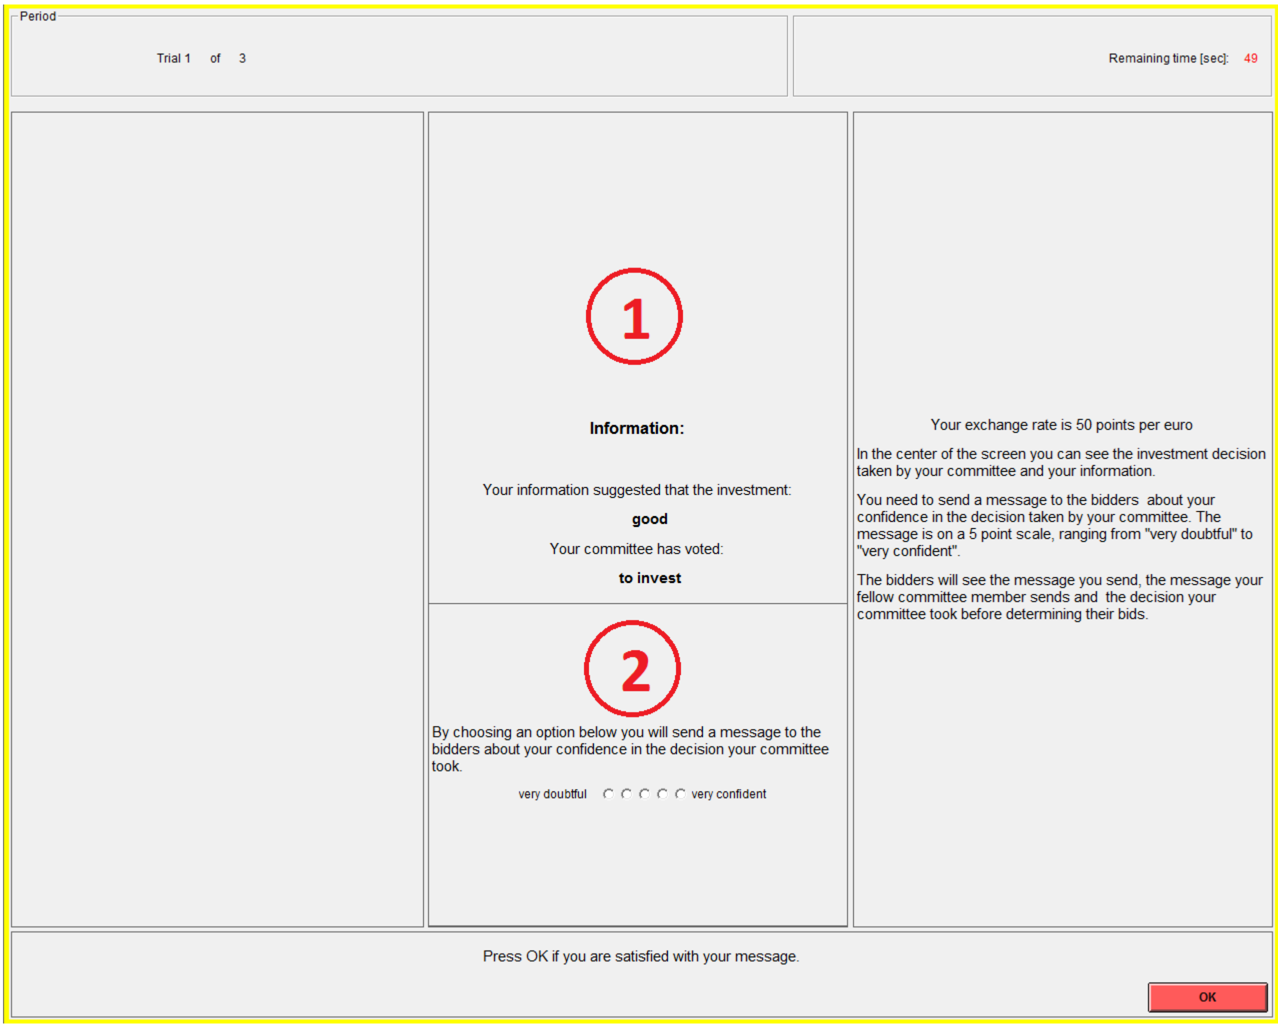

In the center, the top box (indicated with 1) of the screen each committee member is informed about the investment decision taken by his committee, and reminded of what his information suggested. Then he is asked to send a message to the four bidders that are matched to his committee about his confidence in the decision taken by his committee. The message is on a 5 point scale, ranging from “very doubtful” to “very confident”. The message should be entered in the center of the screen in the lower box (indicated with 2).

The message and the decision taken by the committee are then displayed to two pairs of bidders in the next stage of the round. The bidders’ task is to hire committee members to invest for him. The bidders compete in pairs in an auction to determine the wage levels. For each individual committee member the highest bidder is determined and he pays a wage equal to the other bid. The committee members receive the prices paid by both winners of the auction as wage earnings.

At the end of the round each committee member sees how much he earned in the auction and from the investment if one is made.

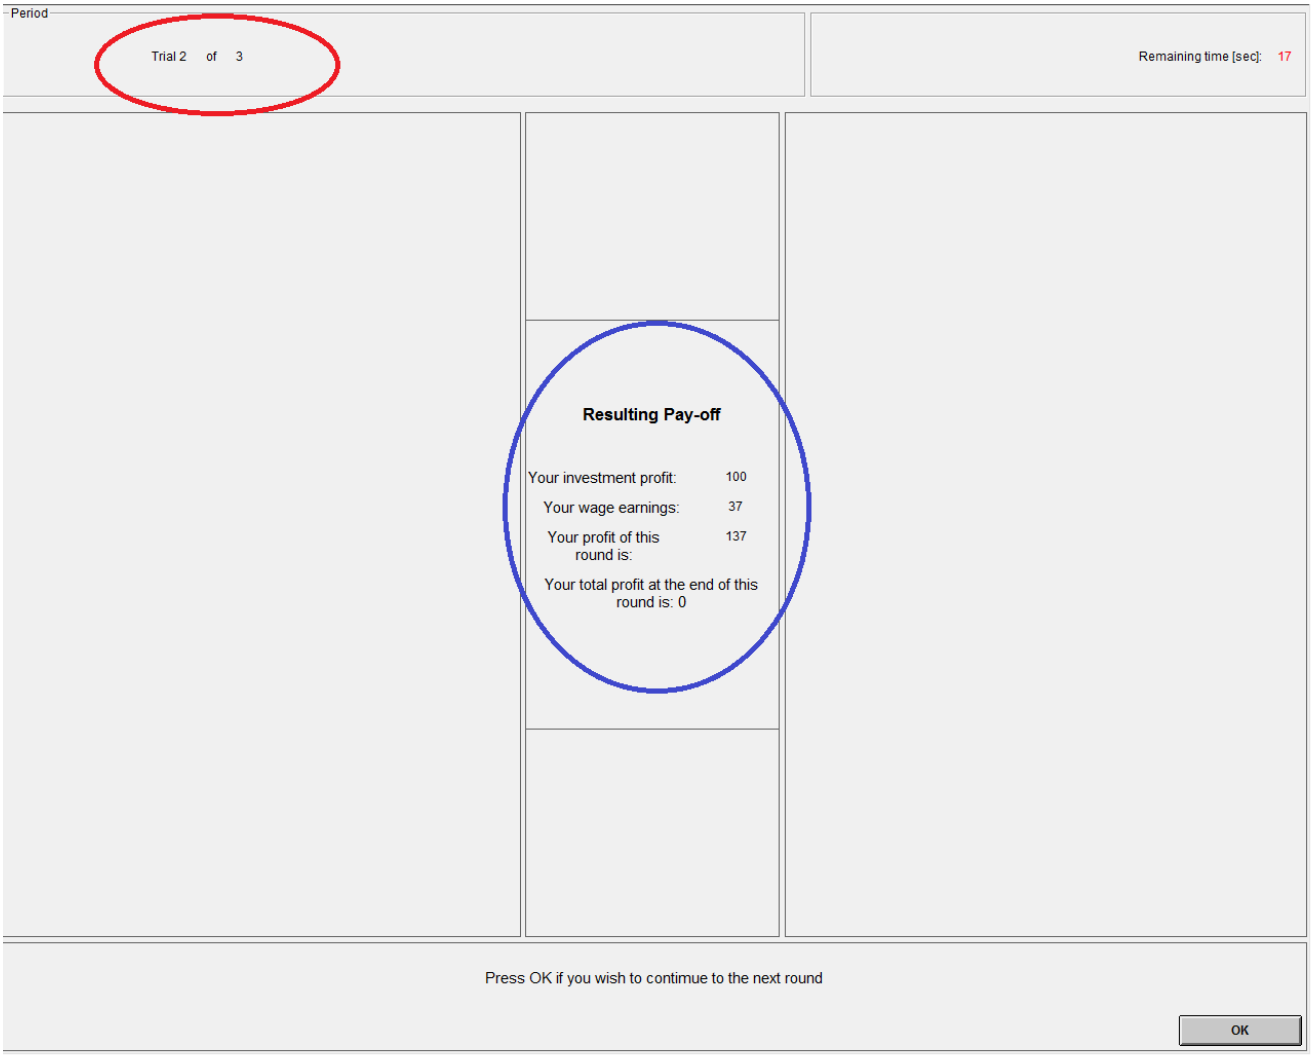

In the center of the screen the total pay-off of the round is displayed (indicated in blue). Both the profit made in this round and the total profit, the sum of all points earned in all rounds up to that point, are displayed. Note that in the example screen the player is still in a trial period. Trial periods do not influence your final pay-off such that the total profit in the example is 0. Whether or not you are in a trial period is always visible in the top of your screen (indicated with a red circle). The wage earnings are determined by two auctions, they are the sum of the payments made by the two winning bidders. The result screen disappears after clicking “OK”, or after 30 seconds automatically and then the next round begins. If you make losses these will be deducted from your show-up fee.

The tasks: bidders

Each bidder can earn money by hiring investors with good information. He makes 4 wage offers each round by submitting bids for the services of the four individual committee members. Each round two committees are randomly matched to two pairs of bidders. Each bidder's screen will display information about the 4 committee members in these 2 committees. One committee is displayed on the left-hand side of the screen and the other committee is displayed at the right-hand side of the screen.

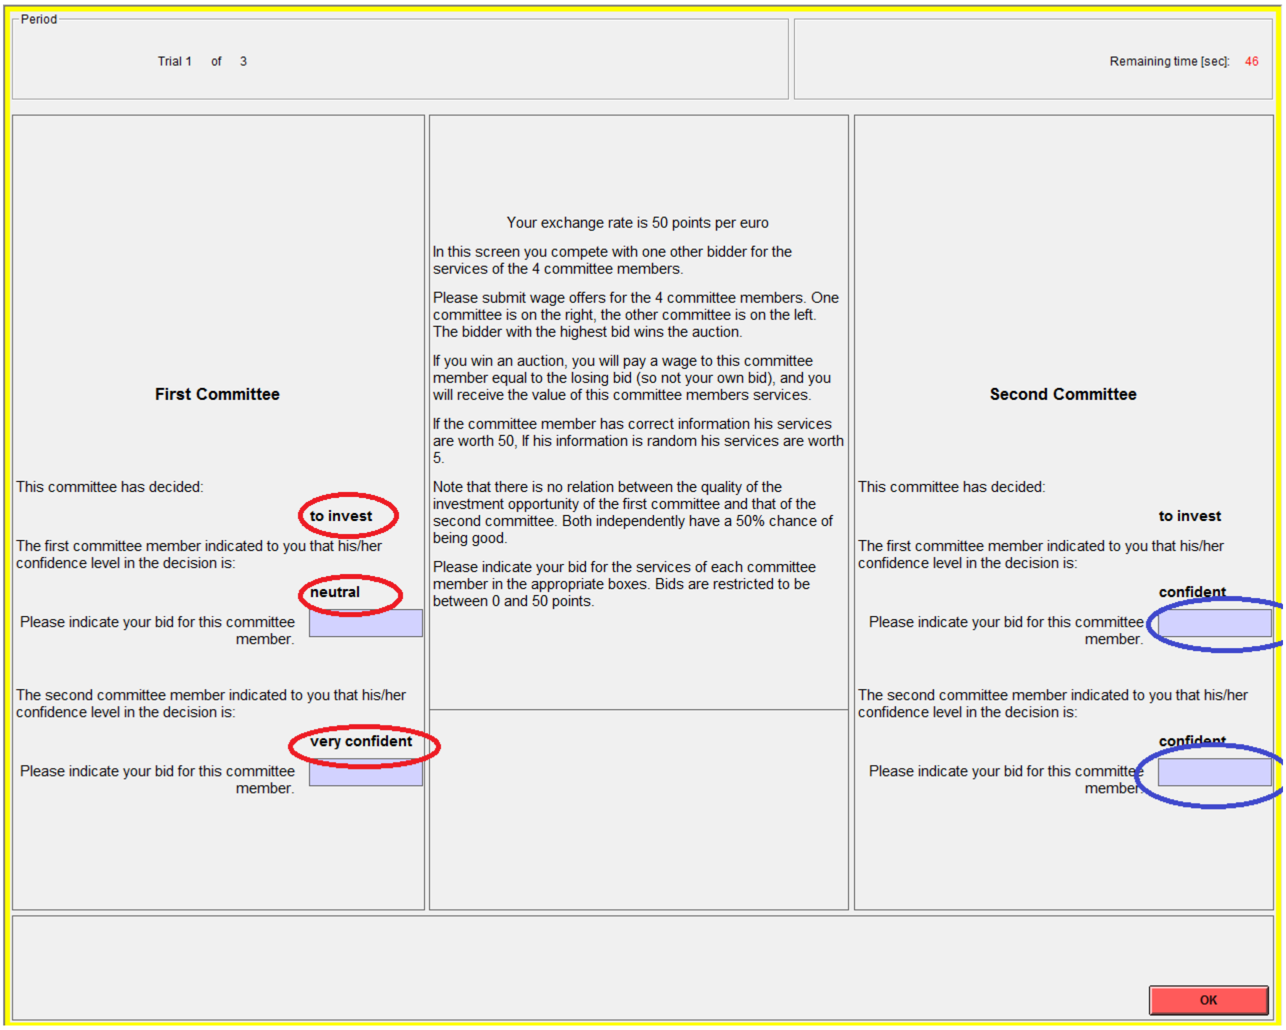

In the screen above each bidder observes whether or not each committee invested and the message sent by each individual committee member (indicated in red only on the left), ranging from "very doubtful" to "very confident". Each bidder then competes in an auction against one other bidder by submitting a bid between 0 and 50 points for each individual committee member. The four bids are submitted by typing them in the appropriate boxes (indicated only on the right in blue) and pressing "OK". The bids are compared for each individual committee member and the highest bidder of the pair wins. In case of a draw the computer randomly determines the winner.

The winner of an auction gets the services of that committee member for this round and pays a wage equal to the loosing bid. The value of the services is 50 points if the committee member had correct information and 5 points if the committee member had random information. Note, the winner does not pay his own bid but a price that is below or equal to his bid, and the minimum value of the services of a bidder is 5 points. Recall: the chance a committee member receives correct information a particular round is 50%. Also note that the quality of the investment opportunities of the two committees is independent of each other.

After each round the bidders see their results in a screen like the one below.

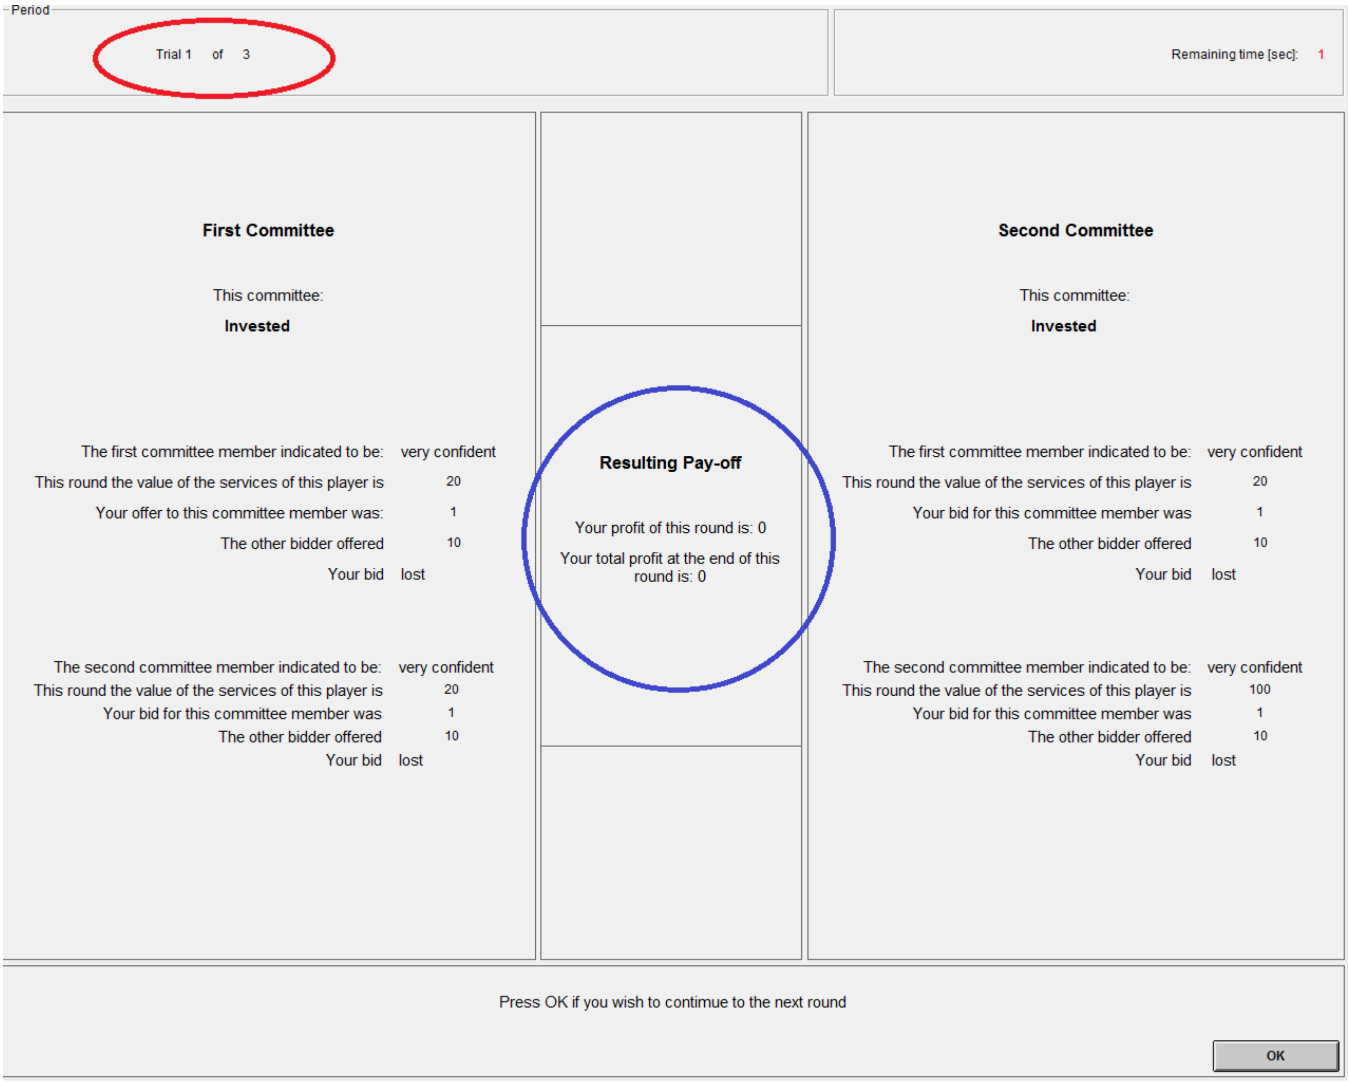

In the center of the screen the pay-off of the round is displayed (indicated in blue). Both the profit made in this round and the total profit, the sum of all points earned in all rounds up to that point, are displayed. Note that in the example screen the player is still in a trial period. Trial periods do not influence your final pay-off such that the total profit in the example is 0. Whether or not you are in a trial period is always visible in the top of your screen (indicated with a red circle). The result screen disappears after clicking “OK” or after 30 seconds automatically and then the next round begins. If you make losses these will be deducted from your show-up fee.

The committees and committee members are displayed in the same order as in the auction screen. So a bidder can easily see which of the auctions he won and which he lost. Note that each round the committees are randomly reordered. Furthermore, whether any particular committee member is presented as first or as second member (top or bottom on the auction and result screen) is also randomly varied per round. It is therefore impossible to track an individual committee member by his position on a bidder’s screen between different rounds. Only for the auction and the result screen in a single round the order is the same.

Summary conclusions

This dissertation deals with some aspects related to gaming, that is, it deals with people that try to take advantage of vulnerabilities and predictabilities in a system.

The first two chapters investigate, in the special case of the tax-system, the requirements on the environment and individuals for gaming to be a problem. Although the first chapter does not present a strict boundary (no necessary conditions are given), a sufficient conditions for the absence of gaming can be provided. Individuals need to have both the opportunity and the desire to deviate from the planners plan. The individuals need to have enough possible choices to be able to deviate from the planner's plan and they need to have a sufficient disagreement with the planner on what is desirable, before they game the system. Although these results seem straightforward, the description of our results shows that most models in the optimal taxation literature satisfy the sufficient conditions for the absence of gaming and that straightforward generalizations of the existing models can be solved. This is both an optimistic and a cautionary tale. On the one hand, it validates much research that has been done and shows how the results of that research can translate into practical outcomes. Our results show that most outcomes from the workhorse model can be translated to a more generalized setting and thus are not particular to the relatively simple models used. On the other hand, however, it shows that much of the literature focuses on a very special model where there are no externalities, no interactions between individuals' utility and the planner has access to a set of incredibly powerful tools. Ironically enough, we can now show that in this special case the planner actually does not need all of these tools, he can do with a slightly less powerful set. What happens outside of this special case, or when the planner has less tools, has not been studied quite as extensively yet, this could provide an interesting avenue for further research.

The next chapter studies the problem of a planner that wants to redistribute from the 'haves' to the 'have-nots'. Such models have traditionally been single dimensional. They are traditionally based on the assumption that differences between individuals can be summarized in a single number, usually assumed to be earnings ability, and individuals make a single dimensional choice, the amount of money they earn. We add to this literature by allowing individuals to differ in many aspects and make even more choices. In this generalized model we derive the Nash-equilibrium using the standard first-order approach. We show how the Nash-equilibrium in this generalized model relates to that of the single dimensional model by deriving generalizations of some of the most well-known results in the literature. We show that the equilibrium forces in the model can still be summarized by the famous ABC-formula: the second best allocation is characterized by the benefits of redistribution, the efficiency cost of the choice distortions and some aspects of the type-distribution. Using this description we can show that the Atkinson-Stiglitz theorem, that shows that indirect taxation is superfluous, is particular to the single-dimensional setting. Only if the planner wants to redistribute in just one dimension, can indirect taxation be superfluous. In general the planner will need as many tax tools as there are dimensions of heterogeneity. In similar vein, it is shown that Mirrlees' (1976) result of separability of the marginal tax rates also does not generalize. If individuals are different in several dimensions and the planner therefore has to use several tax-instruments, these instruments will generally depend on each other. This is the first theoretical result that could explain why many universal health care systems, systems that redistribute from healthy to unhealthy, have payments that depend on the income of the recipient, and are thus tied into the income redistribution. Furthermore, we show that a no-distortion at the top/bottom result holds in this generalized model just as in the uni-dimensional model, that individuals at the low-end of the health and ability distribution might be bunched, as well as link our model to a new strand of the existing literature.

For gaming to be a serious issue, the individuals have to understand their environment and the strategic interactions within it. The last two chapters therefore turn to behavior of the individuals to see how they respond to stimuli and whether our theoretical models can predict their behavior.

The third chapter presents us with something of a puzzle in this regard. The results of our dynamic survey show that experienced players like corporate directors are influenced not just by their environment and their preferences (as assumed by theory), but also by the role they are assigned in the decision process. Although this could be a consequence of informational asymmetries in the normal day-to-day operation of their companies and the response could be appropriate there, this asymmetry does not exist in our survey. The fact that the effect of role persists implies that it is not a calculated rational response to the environment at hand, but comes from somewhere else. The role patterns found, moreover, do not comply with the expectations of the respondents themselves. The respondents indicate they expect the CEO to be most willing to take risks on average. If we ask the respondents to make investment decisions, however, it is the average CFO that is willing to take most risks. The divergence between the expectations of the executives and the choices of these executives creates another puzzle, if they know their environment as well as we assume in theory, this difference should not exist. In the end this chapter leaves the reader as many questions as it answers, but it does show that our approach to studying this hard to reach group is a viable one, so that we can hope to answer them at some later time.

The fourth chapter focuses directly on individuals' strategic interactions. In a series of experiments half of the subjects, our committees of experts, were given private information and the incentive to 'fool' the other half, the reputational market, into thinking they are able experts. The difficult thing was that each expert created several signals and each market participant saw the signals created by 4 experts. Like in the real world, the market has a plethora of signals it can respond to. In theory this multitude of signals should not matter, theory predicts that some of them contain information while others are useless, while both the market and the experts should realize which is which. In our experiments the predictions did not bear out as crisp as in theory. All signals created by the experts seem to contain some information, which means our experts are not maximizing pay-offs. There are also indications that the market is not behaving fully optimal. However, the theory correctly predicts which signals contain most information and are most important to the market. Theory also correctly predicts how the experts will respond to the incentives created by the market, they falsify the signals of ability that are most easily (or cheaply) falsified and try to convince the market of their value. So although there are clear differences between theoretical predictions and observed behavior, the theory is accurate enough to help us identify environments in which career concerns of informed experts can cause problems, without actually having to be an expert in this environment. This implies that we can hope to use theoretical analysis to identify problems created by career concerns, and design procedures that are less susceptible to gaming by the experts, without being an expert in each individual field.

Nederlandse Samenvatting (Summary in Dutch)

Deze dissertatie behandelt aspecten gerelateerd aan het begrip ”gaming”, dit is het gedrag van mensen die misbruik maken van voorspelbare reacties of zwakheden van het systeem om er zelf beter van te worden.

Hoofdstuk twee en drie bestuderen, in het speciale geval van een belastingsysteem, welke eigenschappen van de omgeving en individuen ervoor zorgen dat het systeem gevoelig is voor misbruik. Hoewel het tweede hoofdstuk geen strikte grens aangeeft (we hebben geen noodzakelijke condities kunnen vinden), geeft het wel een voldoende voorwaarde voor de afwezigheid van dit misbruik. Het hoofdstuk laat zien dat individuen zowel de mogelijkheid als de wil moeten hebben om af te wijken, voordat gaming plaatsvindt. Dit betekent dat ze voldoende keuzemogelijkheden moeten hebben om van het plan van de planner af te wijken en het voldoende oneens moeten zijn met de planner, voordat individuen daadwerkelijk misbruik proberen te maken van het systeem. Hoewel deze resultaten vanzelfsprekend lijken, is het bewijs verre van dat. Het bewijs en de beschrijving van de voorwaarden laten zien dat de meeste modellen in de literatuur over optimale belastingsystemen voldoen aan de condities voor de afwezigheid van misbruik. Sterker nog, onze resultaten zijn van toepassing op veel meer modellen dan degenen die gebruikt worden in de literatuur. De resultaten zijn zowel positief als een waarschuwing voor de literatuur. Aan de ene kant laten ze zien hoe de resultaten uit de literatuur, en veel potentiële resultaten uit mogelijke uitbreidingen, omgezet kunnen worden naar een praktisch belastingsysteem. Aan de andere kant laten ze zien dat heel veel van de literatuur focust op een heel bijzondere setting, zonder externaliteiten, zonder interacties in de nutsfuncties van individuen en met een planner met zeer krachtige instrumenten. Ironisch genoeg laten onze resultaten zien dat de planner in deze setting niet al die instrumenten nodig heeft, hij kan met een kleinere set aan instrumenten af. Wat er gebeurt als we niet in deze bijzondere setting zitten, of als de planner een kleinere set aan instrumenten heeft, is een stuk minder bestudeerd en kan een zeer interessante richting zijn voor toekomstig onderzoek.

Het derde hoofdstuk bestudeert het probleem van de centrale planner die wil herverdelen van individuen die het ”goed” hebben naar mensen die het ”slecht” hebben. Zulke modellen bestaan al vrij lang, maar traditioneel zijn ze eendimensionaal: ze zijn gebaseerd op de assumptie dat individuen slechts verschillen in hun vermogen om geld te verdienen en de enkele keuze maken hoeveel geld te verdienen. Onze toevoeging op de bestaande modellen is dat we toestaan dat individuen verschillen in veel eigenschappen en meer keuzes maken. In dit algemenere model laten we zien hoe de meest bekende resultaten uit de eendimensionale literatuur generaliseren naar dit algemener model. In beide modellen wordt de optimale herverdeling bepaald door een drietal krachten: A, het nut dat de planner ontleend aan de herverdeling, B, het verlies dat veroorzaakt wordt door de verstoorde keuzes ten gevolge van de belastingen en C, de verdeling van de types waartussen herverdeling plaatsvindt. Deze beschrijving, die bekend staat als de ABC-formule, is één van de handigste manieren gebleken om het eendimensionale model naar data te brengen. Dat er een directe meerdimensionale versie bestaat is daarom zeer hoopvol voor toekomstig onderzoek. De hoeveelheid eigenschappen waarin individuen verschillen heeft echter ook invloed op de belastingmaatregelen die de planner nodig heeft. In een eendimensionaal model heeft de planner maar één maatregel nodig, de inkomstenbelasting. Dit resultaat (het Atkinson-Stiglitz theorema) generaliseert niet naar meerdere dimensies. Als de planner wil herverdelen in meerdere eigenschappen, dan heeft hij ook meerdere maatregelen nodig. Als we willen herverdelen van gezonde naar minder gezonde mensen, terwijl we tegelijkertijd van mensen met een hoog naar mensen met een laag inkomen herverdelen, dan hebben we daar twee ingrepen voor nodig, bijvoorbeeld een algemene gezondheidsverzekering en de inkomstenbelasting. Deze twee ingrepen moeten ook nog onderling afhankelijk zijn. Die afhankelijkheid zien we wel in de praktijk, maar was theoretisch nog nooit verklaard. Een regeling waarbij de inkomstenbelasting afhankelijk is van de hoeveelheid ziektekosten (bijvoorbeeld door belastingvoordelen voor chronisch zieken zoals in de Nederlandse Wtcg), kan heel goed verklaard worden door een planner die herverdeeld in gezondheid en inkomen. Het hoofdstuk laat ook een paar generalisaties van technische resultaten zien, die de link tussen het eendimensionale en meerdimensionale model duidelijk laten zien en het model ook duidelijker plaatsen in de relatie met andere meerdimensionale 'screening' problemen, zoals de monopolist of veilingmeester die een aantal productlijnen probeert te verkopen.

Voordat misbruik van het systeem een serieus probleem kan zijn moeten de individuen in het systeem hun omgeving en de interacties in deze omgeving voldoende begrijpen. De laatste twee hoofdstukken van deze dissertatie behandelen daarom het gedrag van individuen om te zien hoe ze reageren op prikkels en of onze theoretische modellen hun gedrag kunnen voorspellen.

Het vierde hoofdstuk levert in dit opzicht een aantal vragen op. De survey resultaten laten zien dat bedrijfsdirecteuren niet alleen door hun omgeving en preferenties beÃ¯nvloedt worden, maar ook door hun rol in het besluitvormingsproces. De verschillen in reacties zouden een consequentie kunnen zijn van informatie asymmetrie in de normale gang van zaken. Dat dit roleffect blijft bestaan in onze survey impliceert dat het geen ingecalculeerde respons is op de directe omstandigheden, maar ergens anders vandaan komt. In de survey bestaat geen informatieasymmetrie en is het roleffect dus niet rationeel. De rolpatronen komen ook niet overeen met de verwachtingen van onze respondenten. De respondenten geven aan dat ze, van alle directeuren, de gemiddelde CEO als meest risicotolerant schatten. Als we kijken naar de manier waarop de directeuren investeringsvraagstukken beantwoorden, dan blijkt echter de CFO gemiddeld risicotoleranter te zijn. Dit verschil tussen verwachtingen en uitkomsten is net zo onverwacht als de rolpatronen zelf. Als we er van uitgaan dat de directeuren hun omgeving en de mensen daarin goed kennen, dan zou dit verschil er eigenlijk niet moeten zijn. Met deze bevindingen levert dit hoofdstuk dus net zo veel vragen op als ze beantwoordt. Het onderzoek laat wel zien dat onze onderzoeksaanpak het mogelijk maakt om het gedrag van deze moeilijk te bereiken groep te bestuderen. Er is dus goede hoop dat de vragen in de toekomst beantwoordt kunnen worden.

Het vijfde hoofdstuk focust op strategische interacties tussen individuen, labsubjecten in dit geval. In de experimenten worden de helft van de subjecten aangewezen als experts, zij kregen private informatie en financiële prikkels om de andere helft van de subjecten, de reputatie markt, te overtuigen dat ze goed zijn in het nemen van de relevante besluiten. De grote moeilijkheid in het experiment was, dat elke expert meerdere signalen kon gebruiken om met de markt te communiceren en iedere marktparticipant de signalen van 4 experts zag. Dit geeft de markt, net als in de echte wereld, een plethora aan signalen waar ze op zou kunnen reageren. In theorie zou de hoeveelheid signalen niet uit moeten maken, sommige signalen zijn moeilijk te vervalsen en zouden informatie kunnen bevatten, anderen zijn gemakkelijk te vervalsen en zouden geen informatie moeten bevatten. De theorie voorspelt precies welke signalen geloofd kunnen worden en welke genegeerd moeten worden. Zowel de theoretische markt, als de experts in de theorie snappen dit en dus realiseren ze zich op welke signalen ze zouden moeten focussen. In ons experiment zijn de uitkomsten iets minder scherp dan de theorie. Alle signalen bevatten enige informatie, wat laat zien dat onze experts het spel niet perfect spelen, en er zijn kleine aanwijzingen dat ook de markt niet perfect reageert op de signalen. Als we kijken naar het belang van de signalen, dan blijkt wel dat de signalen die door de theorie aangewezen worden als informatief, meer informatie bevatten dan de theoretisch nutteloze signalen. De markt hecht ook meer waarde aan deze informatieve signalen. De theorie geeft aan dat de experts de simpel te vervalsen signalen gaan vervalsen, als de markt ze positief waardeert en dat is ook precies wat er gebeurd. Dus, ondanks dat er duidelijke afwijkingen zijn van de theoretische voorspellingen, is de theorie accuraat genoeg om situaties te identificeren waar carriÃšre- overwegingen van experts problemen op kunnen leveren, zonder dat we een expert hoeven te zijn in deze specifieke situaties. We kunnen de theoretische analyse dus gebruiken om het besluitvormingsproces zo in te richten dat de mogelijke negatieve gevolgen van carriÃšreoverwegingen beperkt worden, zonder experts te hoeven worden op het inhoudelijke vlak van het besluitvormingsproces.

Bibliography

Adams, William J., and Janet L. Yellen (1976) 'Commodity bundling and the burden of monopoly.' *Quarterly Journal of Economics* 90(3), 475–498

Albanesi, Stefania, and Christopher Sleet (2006) 'Dynamic optimal taxation with private information.' *Review of Economic Studies* 73(1), 1–30

Armstrong, Mark (1996) 'Multiproduct nonlinear pricing.' *Econometrica* 64(1), 51–75

— (2000) 'Optimal multi-object auctions.' *Review of Economic Studies* 67(3), 455–81

Armstrong, Mark, and Jean-Charles Rochet (1999) 'Multi-dimensional screening: A user's guide.' *European Economic Review* 43(4), 959–979

Athey, Susan, and Jonathan Levin (2001) 'Information and competition in US forest service timber auctions.' *Journal of Political Economy* 109(2), 375–417

Atkinson, Anthony B., and Joseph E. Stiglitz (1976) 'The design of tax structure: Direct versus indirect taxation.' *Journal of Public Economics* 6(1-2), 55–75

Baker, Malcolm, and Jeffrey Wurgler (2011) 'Behavioral corporate finance: An updated survey.' Cambridge, MA

Barberis, Nicholas, and Richard Thaler (2003) 'A survey of behavioral finance.' *Handbook of the Economics of Finance* 1, 1053–1128

Ben-David, Itzhak, and John R. Graham (2013) 'Managerial miscalibration.' *Quarterly Journal of Economics* 128(4), 1547–1584

Berg, Joyce E., John W. Dickhaut, and Chandra Kanodia (2009) 'The role of information asymmetry in escalation phenomena: Empirical evidence.' *Journal of Economic Behavior & Organization* 69(2), 135–147

Berger, Allen N., Thomas Kick, and Klaus Schaeck (2013) 'Executive board composition and bank risk taking.' *Journal of Corporate Finance*

Bierbrauer, Felix (2009) 'A note on optimal income taxation, public goods provision and robust mechanism design.' *Journal of Public Economics* 93(5), 667–670

Boadway, Robin, and Bas Jacobs (forthcoming) 'Optimal linear commodity taxation under optimal non-linear income taxation.' *Journal of Public Economics*

Boadway, Robin, and Pierre Pestieau (2011) 'Indirect taxes for redistribution: Should necessity goods be favored?' CESifo Working Paper Series No. 3667 Munich

Bovenberg, A. Lans, and Bas Jacobs (2005) 'Redistribution and education subsidies are Siamese twins.' *Journal of Public Economics* 89(11-12), 2005–2035

Brendon, Charles (2013) 'Efficiency, equity, and optimal income taxation.' Max Weber Program Working Paper MWP 2013/22 Florence

Breusch, Trevor S. (1987) 'Maximum likelihood estimation of random effects models.' *Journal of Econometrics* 36(3), 383–389

Cai, Hongbin, and Joseph Tao-Yi Wang (2006) 'Overcommunication in strategic information transmission games.' *Games and Economic Behavior* 56(1), 7–36

Camerer, Colin, and Dan Lovallo (1999) 'Overconfidence and excess entry: An experimental approach.' 89(1), 306–318

Chakraborty, Indranil (1999) 'Bundling decisions for selling multiple objects.' *Economic Theory* 13(3), 723–733

Chevalier, Judith, and Glenn Ellison (1999) 'Career concerns of mutual fund managers.' *Quarterly Journal of Economics* 114(2), 389–432

Choné, Philippe, and Guy Laroque (2010) 'Negative marginal tax rates and heterogeneity.' *American Economic Review* 100(5), 2532–2547

Christiansen, Vidar (1984) 'Which commodity taxes should supplement the income tax?' *Journal of Public Economics* 24(2), 195–220

Chung, Ji-Woong, Berk A. Sensoy, Léa Stern, and Michael S. Weisbach (2012) 'Pay for performance from future fund flows: The case of private equity.' *Review of financial studies* 25(11), 3259–3304

Clausen, Andrew, and Carlo Strub (2013) 'A general and intuitive evelope theorem.' mimeo Zurich

Clement, Michael B., and Y. Tse Senyo (2005) 'Financial analyst characteristics and herding behavior in forecasting.' *Journal of Finance* 60(1), 307–341

Cremer, Helmuth, Pierre Pestieau, and Jean-Charles Rochet (2001) 'Direct versus indirect taxation: The design of the tax structure revisited.' *International Economic Review* 42(3), 781–800

Dasgupta, Partha S., Peter J. Hammond, and Eric S. Maskin (1979) 'The implementation of social choice rules: Some general results on incentive compatibility.' *Review of Economic Studies* 46(2), 185–216

De Groot, Bert E.A., Sander Renes, René Segers, and Philip-Hans H.B.F. Franses (2012) 'Risk perception and decision-making by the corporate elite: Empirical evidence for netherlands-based companies'

Diamond, Peter A. (1998) 'Optimal income taxation: An example with a u-shaped pattern of optimal marginal tax rates.' *American Economic Review* 88(1), 83–95

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner (2011) 'Individual risk attitudes: Measurement, determinants, and behavioral consequences.' *Journal of the European Economic Association* 9(3), 522–550

Eckel, Catherine C., and Philip J. Grossman (2008) 'Men, women and risk aversion: Experimental evidence.' *Handbook of experimental economics results* 1(7), 1061–73

Erosa, Andres, and Martin Gervais (2002) 'Optimal taxation in life-cycle economies.' *Journal of Economic Theory* 105(2), 338–369

Ewerhart, Christian, and Karsten Fieseler (2003) 'Procurement auctions and unit-price contracts.' *RAND Journal of Economics* 34(3), 569–581

Fama, Eugene F. (1980) 'Agency problems and the theory of the firm.' *Journal of Political Economy* pp. 288–307

Farhi, Emmanuel, and Iván Werning (2010) 'Progressive estate taxation.' *Quarterly Journal of Economics* 125(2), 635–673

— (2013) 'Insurance and taxation over the life cycle.' *Review of Economic Studies* 80(2), 596–635

Favaro, Paul (2001) 'Beyond bean counting: The CFO's expanding role.' *Strategy & Leadership* 29(5), 4–8

Fehr, Ernst, and John A. List (2004) 'The hidden costs and returns of incentives-trust and trustworthiness among CEOs.' *Journal of the European Economic Association* 2(5), 743–771

Fehrler, Sebastian, and Niall Hughes (2014) 'How transparency kills information aggregation: Theory and experiment.' Mimeo Zurich

Fischbacher, Urs (2007) 'Z-tree: Zurich toolbox for ready-made economic experiments.' *Experimental economics* 10(2), 171–178

Fleurbaey, Marc (2006) 'Is commodity taxation unfair?' *Journal of Public Economics* 90(10), 1765–1787

Fleurbaey, Marc, and Francois Maniquet (2006) 'Fair income tax.' *Review of Economic Studies* 73(1), 55–83

Fleurbaey, Marc, and François Maniquet (2007) 'Help the low skilled or let the hardworking thrive? A study of fairness in optimal income taxation.' *Journal of Public Economic Theory* 9(3), 467–500

Gauthier, Stephane, and Guy Laroque (2009) 'Separability and public finance.' *Journal of Public Economics* 93(11), 1168–1174

Gerritsen, Aart (2013) 'Optimal taxation when people do not maximize well-being.' mimeo Erasmus University Rotterdam

Gibbons, R., and K.J. Murphy (1992) 'Optimal incentive contracts in the presence of career concerns: Theory and evidence.' *Journal of Political Economy* 100(3), 468–505

Gillette, Ann B., Thomas H. Noe, and Michael J. Rebello (2008) 'Board structures around the world: An experimental investigation.' *Review of Finance* 12(1), 93–140

Goel, Anand M., and Anjan V. Thakor (2008) 'Overconfidence, CEO selection, and corporate governance.' *Journal of Finance* 63(6), 2737–2784

Goeree, Jacob K., and Leeat Yariv (2011) 'An experimental study of collective deliberation.' *Econometrica* 79(3), 893–921

Golosov, Mikhail, Aleh Tsyvinski, and Iván Werning (2007) 'New dynamic public finance: A user's guide.' In *NBER Macroeconomics Annual 2006, Volume 21,* ed. Daron Acemoglu, Kenneth Rogoff, and Micheal Woodford (Cambridge-MA: MIT Press) pp. 317–388

Golosov, Mikhail, Aleh Tsyvinski, and Matthew Weinzierl (2013) 'Preference heterogeneity and optimal commodity taxation.' *Journal of Public Economics* 97(1), 160–175

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski (2011) 'Optimal dynamic taxes.' NBER Working Paper No. 17642 Cambridge, MA

Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski (2003) 'Optimal indirect and capital taxation.' *Review of Economic Studies* 70(3), 569–587

Graham, John R., Campbell R. Harvey, and Manju Puri (2013) 'Managerial attitudes and corporate actions.' *Journal of Financial Economics* 109(1), 103–121

Graham, John R., Campbell R. Harvey, and Shiva Rajgopal (2005) 'The economic implications of corporate financial reporting.' *Journal of Accounting and Economics* 40(1), 3–73

Greiner, Ben (2004) 'An online recruitment system for economic experiments'

Griffin, Dale, and Amos Tversky (1992) 'The weighing of evidence and the determinants of confidence.' *Cognitive psychology* 24(3), 411–435

Guesnerie, Roger (1995) *A contribution to the pure theory of taxation,* vol. 25 (Cambridge-MA: Cambridge University Press)

Guesnerie, Roger, and Kevin Roberts (1984) 'Effective policy tools and quantity controls.' *Econometrica* 52(1), 59–86

Hambrick, Donald C. (2007) 'Upper echelons theory: An update.' *Academy of Management Review* 32(2), 334–343

Hammond, Peter J. (1979) 'Straightforward individual incentive compatibility in large economies.' *Review of Economic Studies* 46(2), 263–282

_ (1987) 'Markets as constraints: Multilateral incentive compatibility in continuum economies.' *Review of Economic Studies* 54(3), 399–412

Hellwig, Martin F. (2010) 'A generalization of the Atkinson–Stiglitz (1976) theorem on the undesirability of nonuniform excise taxation.' *Economics Letters* 108(2), 156–158

Hermalin, Benjamin E., and Michael S. Weisbach (2014) 'Understanding corporate governance through learning models of managerial competence.' *NBER Working Paper No. w20028*

Holmström, Bengt (1999) 'Managerial incentive problems: A dynamic perspective.' *Review of Economic Studies* 66(1), 169–182

Holmstrom, Bengt, and Paul Milgrom (1987) 'Aggregation and linearity in the provision of intertemporal incentives.' *Econometrica* pp. 303–328

Hong, Harrison, Jeffrey D. Kubik, and Amit Solomon (2000) 'Security analysts' career concerns and herding of earnings forecasts.' *RAND Journal of Economics* 31(1), 121–144

Hurwicz, Leonid (1960) 'Optimality and informational efficiency in resource allocation processes.' In *Mathematical Methods in the Social Sciences,* ed. Kenneth Arrow, S. Karlin, and P. Suppes (Redwood City, CA: Stanford University Press)

— (1973) 'The design of mechanisms for resource allocation.' *American Economic Review* 63(2), 1–30

Huse, Morten (2005) 'Accountability and creating accountability: A framework for exploring behavioural perspectives of corporate governance.' *British Journal of Management* 16(s1), S65–S79

Jacobs, Bas, and Ruud A. De Mooij (2011) 'Pigou meets Mirrlees: On the irrelevance of tax distortions for the second-best pigouvian tax.' CESifo Working Paper No. 3250 Munich

Jacquet, Laurence, and Dirk Van de Gaer (2011) 'A comparison of optimal tax policies when compensation or responsibility matter.' *Journal of Public Economics* 95(11), 1248–1262

Jacquet, Laurence, and Etienne Lehmann (2014) 'Optimal nonlinear income taxation with multidimensional types: The case with heterogeneous behavioral responses.' THEMA, Université de Cergy-Pontoise Paris

Jacquet, Laurence, Etienne Lehmann, and Bruno Van der Linden (2013) 'Optimal redistributive taxation with both extensive and intensive responses.' *Journal of Economic Theory* 148(5), 1770–1805

Jehiel, Philippe, Moritz Meyer-Ter-Vehn, and Benny Moldovanu (2007) 'Mixed bundling auctions.' *Journal of Economic Theory* 134(1), 494–512

Jensen, Michael, and Edward J. Zajac (2004) 'Corporate elites and corporate strategy: How demographic preferences and structural position shape the scope of the firm.' *Strategic Management Journal* 25(6), 507–524

Judd, Kenneth, and Che-Lin Su (2006) 'Optimal income taxation with multidimensional taxpayer types.' *Computing in Economics and Finance*

Kanbur, Ravi, Matti Tuomala, and Jukka Pirttilae (2006) 'Non-welfarist optimal taxation and behavioral economics.' *Journal of Economic Surveys* 20(5), 1467–6419

Kaplow, Louis (2006) 'On the undesirability of commodity taxation even when income taxation is not optimal.' *Journal of Public Economics* 90(6), 1235–1250

Kaplow, Louis, and Steven Shavell (2001) 'Any non-welfarist method of policy assessment violates the Pareto principle.' *Journal of Political Economy* 109(2), 281–286

Kleven, Henrik J., Claus T. Kreiner, and Emmanuel Saez (2009) 'The optimal income taxation of couples.' *Econometrica* 77(2), 537–560

Koch, Alexander K., Albrecht Morgenstern, and Philippe Raab (2009) 'Career concerns incentives: An experimental test.' *Journal of Economic Behavior & Organization* 72(1), 571–588

Kocherlakota, Narayana R. (2005) 'Zero expected wealth taxes: A mirrlees approach to dynamic optimal taxation.' *Econometrica* 73(5), 1587–1621

_ (2010) *The new dynamic public finance* (Princeton-NJ: Princeton University Press)

Kopczuk, Wojciech (2013) 'Incentive effects of inheritances and optimal estate taxation.' NBER Working Paper No. 18747 Cambridge, MA

Laroque, Guy R. (2005) 'Indirect taxation is superfluous under separability and taste homogeneity: A simple proof.' *Economics Letters* 87(1), 141–144

Lawrence, Barbara S. (1997) 'Perspective-the black box of organizational demography.' *Organization Science* 8(1), 1–22

Leblanc, Richard, and Mark S. Schwartz (2007) 'The black box of board process: Gaining access to a difficult subject.' *Corporate Governance: An International Review* 15(5), 843–851

Lim, Jongha, Berk A. Sensoy, and Michael S. Weisbach (2013) 'Indirect incentives of hedge fund managers.' NBER Working Paper No. 18903

List, John A., and Charles F. Mason (2011) 'Are CEOs expected utility maximizers?' *Journal of Econometrics* 162(1), 114–123

Lockwood, Benjamin B., and Matthew C. Weinzierl (2012) 'De gustibus non est taxandum: Theory and evidence on preference heterogeneity and redistribution.' NBER Working Paper No. 17784 Cambridge, MA

Lorsch, Jay W. (1995) 'Empowering the board.' *Harvard Business Review* 73(1), 107–117

Lorsch, Jay W., and Jack Young (1990) 'Pawns or potentates: The reality of America's corporate boards.' *The Executive* 4(4), 85–87

MacCrimmon, Kenneth R., and Donald A. Wehrung (1990) 'Characteristics of risk taking executives.' *Management Science* 36(4), 422–435

Mace, Myles L. (1972) 'The president and the board of directors.' *Harvard Business Review* 50(12), 37–49

Malmendier, Ulrike, and Geoffrey Tate (2005a) 'CEO overconfidence and corporate investment.' *Journal of Finance* 60(6), 2661–2700

— (2005b) 'Does overconfidence affect corporate investment? CEO overconfidence measures revisited.' *European Financial Management* 11(5), 649–659

— (2008) 'Who makes acquisitions? CEO overconfidence and the market's reaction.' *Journal of Financial Economics* 89(1), 20–43

Mas-Collel, Andreau, Micheal D. Whinston, and Jerry R. Green (1995) *Microeconomics* (Oxford: Oxford University Press)

Maskin, Eric (1999) 'Nash equilibrium and welfare optimality.' *Review of Economic Studies* 66(1), 23–38

McAfee, R. Preston, and John McMillan (1988) 'Multidimensional incentive compatibility and mechanism design.' *Journal of Economic Theory* 46(2), 335–354

Milbourn, Todd T., Richard L. Shockley, and Anjan V. Thakor (2001) 'Managerial career concerns and investments in information.' *RAND Journal of Economics* 32(2), 334–51

Mirrlees, James A. (1971) 'An exploration in the theory of optimum income taxation.' *Review of Economic Studies* 38(2), 175–208

— (1976) 'Optimal tax theory: A synthesis.' *Journal of Public Economics* 6(4), 327–358

Myerson, Roger B. (1979) 'Incentive compatibility and the bargaining problem.' *Econometrica* 47(1), 61–73

Ooghe, Erwin, and Andreas Peichl (2011) 'Fair and efficient taxation under partial control: Theory and evidence.' CESifo Working Paper No. 3518 Munich

Ottaviani, Marco, and Peter Sørensen (2001) 'Information aggregation in debate: Who should speak first?' *Journal of Public Economics* 81(3), 393–421

Palfrey, Thomas R. (1983) 'Bundling decisions by a multiproduct monopolist with incomplete information.' *Econometrica* 51(2), 463–83

Pass, Brendan (2012) 'Convexity and multi-dimensional screening for spaces with different dimensions.' *Journal of Economic Theory* 147(6), 2399–2418

Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2011) 'Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure.' Center for Mathematical Studies in Economics and Management Science Working Paper No. 1501

Piketty, Thomas (2014) *Capital in the 21st Century* (Cambridge: Harvard University Press)

Rabe-Hesketh, Sophia, Anders Skrondal, and Andrew Pickles (2004) 'Generalized multilevel structural equation modeling.' *Psychometrika* 69(2), 167–190

— (2005) 'Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects.' *Journal of Econometrics* 128(2), 301–323

Renes, Sander (2011) 'Balancing the bids, solutions for unit price auctions.' Tinbergen Institute Discussion Paper 11-047/1 Amsterdam

Renes, Sander, and Floris T. Zoutman (2014a) 'As easy as ABC? Multi-dimensional screening in public finance.' Rotterdam. mimeo

— (2014b) 'When a price is enough: Implementation in optimal tax design.' Rotterdam. mimeo

Roberts, Kevin (1984) 'The theoretical limits to redistribution.' *Review of Economic Studies* 51(2), 177–195

— (2000) 'A reconsideration of the optimal income tax.' In 'Incentives and Organization, Papers in Honour of Sir James Mirrlees' (Oxford: Department of Economics (University of Oxford))

Rochet, Jean-Charles (1985) 'The taxation principle and multi-time Hamilton-Jacobi equations.' *Journal of Mathematical Economics* 14(2), 113–128

Rochet, Jean-Charles, and Philippe Choné (1998) 'Ironing, sweeping, and multidimensional screening.' *Econometrica* 66(4), 783–826

Rothschild, Casey, and Florian Scheuer (2011) 'Optimal taxation with rent-seeking.' NBER Working Paper No. 17035 Cambridge, MA

— (2013) 'Redistributive taxation in the Roy model.' *Quarterly Journal of Economics* 128(2), 623–668

— (2014) 'A theory of income taxation under multidimensional skill heterogeneity.' NBER Working Paper No. 19822 Cambridge, MA

Sadka, Efraim (1976) 'On income distribution, incentive effects and optimal income taxation.' *Review of Economic Studies* 43(2), 261–267

Saez, Emmanuel (2001) 'Using elasticities to derive optimal income tax rates.' *Review of Economic Studies* 68(1), 205–229

— (2002a) 'The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes.' *Journal of Public Economics* 83(2), 217–230

— (2002b) ‘Optimal income transfer programs: Intensive versus extensive labor supply responses.’ *Quarterly Journal of Economics* 117(3), 1039–1073

Sanders, W. Gerard (2001) ‘Behavioral responses of CEOs to stock ownership and stock option pay.’ *Academy of Management Journal* 44(3), 477–492

Sandmo, Agnar (1993) ‘Optimal redistribution when tastes differ.’ *FinanzArchiv/Public Finance Analysis* 50(2), 149–163

Scharfstein, David S., and Jeremy C. Stein (1990) ‘Herd behavior and investment.’ *American Economic Review* pp. 465–479

Seade, Jesus K. (1977) ‘On the shape of optimal tax schedules.’ *Journal of Public Economics* 7(2), 203–235

Shavell, Steven (1979) ‘Risk sharing and incentives in the principal and agent relationship.’ *Bell Journal of Economics* 10(1), 55–73

Stiglitz, Joseph E. (1987a) ‘The design of labor contracts: Economics of incentives and risk sharing.’ In *Incentives, cooperation and risk sharing,* ed. H. Nalbantian (Totowa, NJ: Rowman and Littlefield) pp. 47–68

— (1987b) ‘Pareto efficient and optimal taxation and the new new welfare economics.’ *Handbook of Public Economics* 2, 991–1042

Suurmond, Guido, Otto H. Swank, and Bauke Visser (2004) ‘On the bad reputation of reputational concerns.’ *Journal of Public Economics* 88(12), 2817–2838

Swank, Job, Otto H. Swank, and Bauke Visser (2008) ‘How committees of experts interact with the outside world: some theory, and evidence from the fomc.’ *Journal of the European Economic Association* 6(2-3), 478–486

Swank, Otto H, and Bauke Visser (2013) ‘Is transparency to no avail?’ *Scandinavian Journal of Economics* 115(4), 967–994

Trahan, Emery A., and Lawrence J. Gitman (1995) ‘Bridging the theory-practice gap in corporate finance: A survey of chief financial officers.’ *Quarterly Review of Economics and Finance* 35(1), 73–87

Vickrey, William (1961) ‘Counterspeculation, auctions, and competitive sealed tenders.’ *Journal of Finance* 16(1), 8–37

Visser, Bauke, and Otto H. Swank (2007) ‘On committees of experts.’ *Quarterly Journal of Economics* 122(1), 337–372

Werning, Ivan (2007) 'Pareto efficient income taxation.' Mimeo MIT Cambridge, MA

Wiersema, Margarethe F., and Karen A. Bantel (1992) 'Top management team demography and corporate strategic change.' *Academy of Management Journal* 35(1), 91–121

Wilson, Robert (1996) 'Nonlinear pricing and mechanism design.' vol. 1 (Elsevier) pp. 253–293

Zheng, Charles Z. (2000) 'Optimal auction in a multidimensional world.' Center for Mathematical Studies in Economics and Management Science, Northwestern University, Discussion Paper 1282

Zorn, Dirk M. (2004) 'Here a chief, there a chief: The rise of the CFO in the american firm.' *American Sociological Review* 69(3), 345–364

Zuckerman, Marvin, Elizabeth A. Kolin, Leah Price, and Ina. Zoob (1964) 'Development of a sensation-seeking scale.' *Journal of Consulting Psychology* 28(6), 477–482

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus University Rotterdam, University of Amsterdam and VU University Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. The following books recently appeared in the Tinbergen Institute Research Series:

541. T.WANG, *Essays on Empirical Market Microstructure.*
542. T. BAO, *Experiments on Heterogeneous Expectations and Switching Behavior.*
543. S.D. LANSDORP, *On Risks and Opportunities in Financial Markets.*
544. N. MOES, *Cooperative decision making in river water allocation problems.*
545. P. STAKENAS, *Fractional integration and cointegration in financial time series.*
546. M. SCHARTH, *Essays on Monte Carlo Methods for State Space Models.*
547. J. ZENHORST, *Macroeconomic Perspectives on the Equity Premium Puzzle.*
548. B. PELLOUX, *the Role of Emotions and Social Ties in Public On Good Games: Behavioral and Neuroeconomic Studies.*
549. N. YANG, *Markov-Perfect Industry Dynamics: Theory, Computation, and Applications.*
550. R.R. VAN VELDHUIZEN, *Essays in Experimental Economics.*
551. X. ZHANG, *Modeling Time Variation in Systemic Risk.*
552. H.R.A. KOSTER, *The internal structure of cities: the economics of agglomeration, amenities and accessibility.*
553. S.P.T. GROOT, *Agglomeration, globalization and regional labor markets: Micro evidence for the Netherlands.*
554. J.L. MÖHLMANN, *Globalization and Productivity Micro-Evidence on Heterogeneous Firms, Workers and Products.*
555. S.M. HOOGENDOORN, *Diversity and Team Performance: A Series of Field Experiments.*
556. C.L. BEHRENS, *Product differentiation in aviation passenger markets: The impact of demand heterogeneity on competition.*
557. G. SMRKOLJ, *Dynamic Models of Research and Development.*
558. S. PEER, *The economics of trip scheduling, travel time variability and traffic information.*
559. V. SPINU, *Nonadditive Beliefs: From Measurement to Extensions.*
560. S.P. KASTORYANO, *Essays in Applied Dynamic Microeconometrics.*
561. M. VAN DUIJN, *Location, choice, cultural heritage and house prices.*
562. T. SALIMANS, *Essays in Likelihood-Based Computational Econometrics.*
563. P. SUN, *Tail Risk of Equidity Returns.*
564. C.G.J. KARSTEN, *The Law and Finance of M&A Contracts.*

565. C. OZGEN, *Impacts of Immigration and Cultural Diversity on Innovation and Economic Growth.*

566. R.S. SCHOLTE, *The interplay between early-life conditions, major events and health later in life.*

567. B.N.KRAMER, *Why don't they take a card? Essays on the demand for micro health insurance.*

568. M. KILIÇ, *Fundamental Insights in Power Futures Prices.*

569. A.G.B. DE VRIES, *Venture Capital: Relations with the Economy and Intellectual Property.*

570. E.M.F. VAN DEN BROEK, *Keeping up Appearances.*

571. K.T. MOORE, *A Tale of Risk: Essays on Financial Extremes.*

572. F.T. ZOUTMAN, *A Symphony of Redistributive Instruments.*

573. M.J. GERRITSE, *Policy Competition and the Spatial Economy.*

574. A. OPSCHOOR, *Understanding Financial Market Volatility.*

575. R.R. VAN LOON, *Tourism and the Economic Valuation of Cultural Heritage.*

576. I.L. LYUBIMOV, *Essays on Political Economy and Economic Development.*

577. A.A.F. GERRITSEN, *Essays in Optimal Government Policy.*

578. M.L. SCHOLTUS, *The Impact of High-Frequency Trading on Financial Markets.*

579. E. RAVIV, *Forecasting Financial and Macroeconomic Variables: Shrinkage, Dimension reduction, and Aggregation.*

580. J. TICHEM, *Altruism, Conformism, and Incentives in the Workplace.*

581. E.S. HENDRIKS, *Essays in Law and Economics.*

582. X. SHEN, *Essays on Empirical Asset Pricing.*

583. L.T. GATAREK, *Econometric Contributions to Financial Trading, Hedging and Risk Measurement.*

584. X. LI, *Temporary Price Deviation, Limited Attention and Information Acquisition in the Stock Market.*

585. Y. DAI, *Efficiency in Corporate Takeovers.*

586. S.L. VAN DER STER, *Approximate feasibility in real-time scheduling: Speeding up in order to meet deadlines.*

587. A. SELIM, *An Examination of Uncertainty from a Psychological and Economic Viewpoint.*

588. B.Z. YUESHEN, *Frictions in Modern Financial Markets and the Implications for Market Quality.*

589. D. VAN DOLDER, *Game Shows, Gambles, and Economic Behavior.*

590. S.P. CEYHAN, *Essays on Bayesian Analysis of Time Varying Economic Patterns.*