tinbergen *institute*

Do ut Des

Incentives, Reciprocity, and Organizational Performance

Arjan Non

Research Series

Erasmus Universiteit Rotterdam

Do ut Des: Incentives, Reciprocity, and Organizational Performance

ISBN: 978 90 3610 255 1

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

This book is no. 510 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.

Do ut Des: Incentives, Reciprocity, and Organizational Performance

Do ut Des: Prikkels, reciprociteit, en prestaties van organisaties

Proefschrift

ter verkrijging van de graad van doctor aan de

Erasmus Universiteit Rotterdam

op gezag van de

rector magnificus

Prof.dr. H.G. Schmidt

en volgens het besluit van het College voor Promoties

De openbare verdediging zal plaatsvinden op

donderdag 6 oktober 2011 om 13:30 uur

door

Jan Adrianus Non

geboren te Ridderkerk

ERASMUS UNIVERSITEIT ROTTERDAM

Promotiecommissie

Promotor: Prof.dr. A.J. Dur

Overige leden: Dr. B. Donkers
Prof.dr. H. Oosterbeek
Prof.dr. O.H. Swank

Preface

This thesis is ascribed to one author, but this fails to do justice to several others who played a role in the process of writing this dissertation. Writing a PhD thesis is usually not a completely smooth proces, and this thesis is no exception. It is fair to say that this thesis would not be there without the continuous support of several others. I would like to take the opportunity to thank many of them.

First and foremost, I would like to thank my promotor Robert Dur. Robert, your great enthusiasm and amazing creativity in doing research is highly inspiring. You have the rare quality of turning problems into challenges. Every time I left your office I was full of new ideas. Moreover, I learned a lot from your comments on my work and from working together. I am grateful that your door was always open to discuss our research. Apart from your quality as an outstanding and stimulating researcher, you are also a pleasant person to be around during dinners, lunch breaks, and drinks. You typically show a well-developed sense of humor.

Second, I would like to thank my other co-authors, Josse Delfgaauw, Hein Roelfsema and Willem Verbeke. Cooperating was a highly valuable experience. Josse, you turned out to be a worthy substitute for Robert from time to time. I would also like to seize the opportunity to thank the management of the anonymous retail chain who allowed us to conduct two field experiments in their organisation, and those who were involved in communicating and designing the experiments.

Third, I thank Bas Donkers, Hessel Oosterbeek, and Otto Swank for being part of my doctoral committee and for their helpful comments. Also, I would like to thank the members of the larger committee, Randolph Sloof and Thomas Dohmen. Thomas, I am looking forward to a fruitful cooperation in the coming years at Maastricht University.

Fourth, I would like to thank the many colleagues who played a more indirect role in writing this thesis. Writing a thesis is often thought of as an isolated and lonely endeavour, but the daily interactions with several others prevented that from coming true. The first person to thank is Sjoerd van den Hauwe, with whom I shared an office for 3 years. I also thank Phillippe Versijp, who shared the office with me in the first year of my PhD. Philippe, you made complaining about working conditions, Avenance-food and traffic jams a real pleasure.

Sharing a room clearly stimulates social interaction, the same is true for sharing a promotor. Joeri, Margaretha, Oke, Robin, and Jan, I appreciated the numerous discussions about research, teaching and miscellaneous topics. Joeri and Margaretha, I also enjoyed your company during several conferences, for instance in Granada, Venice and Barcelona. Social interaction is obviously not restricted to people working on similar topics: I have the pleasant memory of several nice chats with Riemer, Paul, Sven, Floris, and many others. Furthermore, I liked participating in the traditional TI-football team, although absolute gallery-play remained far out-of-reach.

I am also happy that many friends made me forget about work altogether. Ben, Hugo, and Kees-Jan, I greatly appreciated your non-economic perspective in discussions on a diversity of topics. Hans, I hope to continue our yearly tradition of hiring an exciting car for a nice trip. Furhermore, I thank Goalenbeer, De Trein, De Muur, De Tornado, Pollo, Abe, The Iceman, Lucky Léon, Killer Miller and Phenomenon for exciting football matches every week, and for the related as well as the unrelated social activities.

This is also an excellent opportunity to thank my family. Mariëlle, by pursuing a PhD as my elder sister you at least made me aware of what PhD-study involves, and your experiences greatly facilitated my life as a PhD. Wilco, it seems that being exposed to such an academic environment is paying off, albeit in a different discipline and, in line with your outgoing character, in a different country from time to time. Have you ever thought about doing a PhD? Of course, I also thank my parents for their unconditional love and support. Despite numerous stays in hotels during conferences and holidays, I never found a hotel that outperformed 'hotel mama'. Finally, I am grateful to God for taking part in His wonderful world.

Contents

Chapter 1

Introduction

"My father taught me to work; he did not teach me to love it. I never did like to work, and I don't deny it. I'd rather read, tell stories, crack jokes, talk, laugh - anything but work." *Abraham Lincoln*

1.1 Motivation

Economists have traditionally viewed the employer-employee relation as a typical example of an agency relationship. The worker (agent) acts on behalf of the employer (principal), and is supposed to advance the employer's interest. However, as the quote above makes clear, the employee's interests may differ considerably from that of the employer. Moreover, in most employment situations, workers have considerable room for pursuing their own goals at the expense of the company. The typical solution economists come up with, is to align the employee's interests with those of the employer using financial incentives. When a worker's remuneration is closely tied to his performance, he has a clear incentive to exert effort on the job. The importance of financial incentives is confirmed in many empirical studies (see Prendergast 1999 and Lazear and Oyer 2009 for an overview). For instance, as early as 1850, Edwin Chadwick managed to increase survival rates among British criminals transported to Australia from 40% to 98% by paying captains for each live convict that disembarked in Australia, instead of the number of prisoners that embarked in England. More recently, Shearer (2004) reports an increase in productivity of about

20% when tree-planters work under piece rates rather than fixed wages. Lazear (2000) finds performance effects of piece rates in the same order of magnitude.

However, few economists will maintain that firms rely exclusively on monetary incentives to motivate workers. The large majority of workers face only weak monetary incentives, despite its proven effectiveness. In a typical Western economy, less than half of all workers receive payments that are explicitly related to their performance, and the median magnitude of this performance component is less than 5% of their base salary, see Piekkola (2005), Lemieux et al. (2007) and Belfield and Marsden (2003). Moreover, workers and managers themselves often do not consider monetary incentives as the most important motivator, see the surveys by Minkler (2004) and Campbell and Kamlani (2007).

In line with these observations, economists have recently strongly qualified the view of workers as untrustworthy work-avoiders and managers as passive contract writers. Economists now widely acknowledge that workers are not exclusively interested in maximizing their own material welfare, and that several other motives play an important role in the workplace. For example, recent papers take into account that workers are to some extent intrinsically motivated (e.g. Delfgaauw and Dur 2007), that workers' effort decisions are influenced by social norms (Sliwka 2007), that workers respond in kind to friendly and hostile acts (Falk and Fischbacher 2006), and that workers desire the manager's respect (Ellingsen and Johannesson 2007). As a further illustration of economists' efforts to depart from the neoclassical homo economicus, what to think of recent titles like "Pride and Prejudice, the Human Side of Incentive Theory" (Ellingsen and Johannesson 2008), "Competition and Incentives with Motivated Agents" (Besley and Ghatak 2005), and "Incorporating Morale into a Classical Agency Model" (Stowe 2009)? Clearly, economists now widely recognize that there are many factors besides material incentives that strongly influence work motivation. Importantly, these motivations may interact with financial incentives in surprising ways. For instance, material incentives sometimes crowd-out intrinsic motivation (Frey and Jegen, 2001), or serve as a signal of a social norm (Sliwka 2007).

In the light of these developments it is perhaps unsurprising that economists have also resorted to new methods of research. Economists now make more and more frequent use of controlled experiments to tease out the several motives that play a role in human behavior. Experiments are often more suitable for this purpose than naturally occurring data, as establishing causality is easier in an environment largely controlled by the experimenter. Thus, much empirical research on workplace behavior now takes the form of a laboratory or field-experiment.

This thesis contributes to this emerging literature. In the first part of this thesis, consisting of chapters 2-3, my co-authors and I analyze agency problems when workers are reciprocal towards the manager's kindness. The focus of the analysis is on optimal wage and incentive provision. In the second part of this thesis I present two field experiments conducted in a Dutch retail chain. Both experiments are designed to test some specific elements of standard tournament theory.

The remainder of this introduction is structured as follows. I will first briefly discuss the literature on reciprocity and the related notions of gift-exchange and conditional altruism. Next, I will outline the contribution of each chapter to the existing literature. Finally, I provide an overview of the contents of each chapter.

1.2 Reciprocity

Economists' strong emphasis on employees' pecuniary motives has long distracted attention away from many other considerations that affect economic outcomes. More recent literature, however, departs from the view that workers are primarily motivated by money, to the exclusion of various other motives. As discussed above, several other considerations have been incorporated in theoretical models or made subject of empirical investigation. The first part of this thesis focuses on one of these motives, namely reciprocity. Reciprocity roughly means that people respond in kind to friendly and hostile acts. The tendency to promote the welfare of a kind person is usually referred to as 'positive reciprocity', whereas the tendency to reduce the welfare of an unkind person is referred to as 'negative reciprocity' (Dohmen et al. 2009).

The importance of reciprocity in the workplace has been brought under the attention of economists by Akerlof in 1982. In his seminal paper titled 'Labor contracts as partial gift-exchange', he argues that the employer's benevolent treatment of workers and worker's effort can be seen as reciprocal gifts. Employees put in more effort than required, while employers treat their employees better than strictly necessary. Thus, labor contracts can, at least partially, be considered as an exchange of gifts, induced by the norm of reciprocity.

The concepts of gift-exchange and reciprocity date back long before Akerlof introduced these concepts into economics. In social psychology for example, Homans wrote in 1958 that "Social behavior is an exchange of goods, material goods but also non-material ones, such as the symbols of approval or prestige. Persons that give much to others try to get much from them, and persons that get much from others are under pressure to give much to them. This process of influence tends to work out at equilibrium to a balance in the exchanges" (p. 606). In sociology, Blau (1964) distinguishes between economic exchange and social exchange. Social exchange is characterised by unspecified obligations. While there is a general expectation that a favor will be returned, the nature of the return is not specified in advance and cannot be bargained over. According to Blau (1964), the reason that favors are reciprocated is that people are "interested in maintaining a balance between inputs and outputs and staying out of debt in their social transactions" (p. 26). A gift induces an obligation, and there is a moral norm to reciprocate. Receiving a gift is therefore a doubtful blessing, which is apparent in the etymology of the word 'gift'. In ancient German and Greek, the same word is used for gift and poison.

Economists now widely recognise that reciprocal motivations may have a substantial impact on wages and the optimality of labor contracts. Numerous laboratory experiments have provided support for Akerlof's hypothesis that paying high wages triggers high effort (Fehr and Gächter 1998), although recent field experiments are less supportive (e.g. Gneezy and List 2006, List 2006, Al-Ubaydli et al. 2007, Hennig-Schmidt et al. 2008, Bellemare and Shearer 2009). Moreover, a field experiment conducted by Kube et al. (2006) shows that a gift in-kind results in a substantial increase in workers' productivity. Furthermore, in several laboratory experiments, generous implicit contracts are more succesful in eliciting effort than contracts that explicitly specify fines in case of verified shirking (Fehr and Gächter 2000). In line with these experiments, survey evidence shows that individuals who are more inclined to reciprocate an employer's kindness earn more, are less likely to be unemployed, and work more overtime, especially when they perceive their wage as fair (Dohmen et al. 2009). These findings suggest that employers benefit from positive reciprocity, at least when they treat workers appropriately. On the other hand, employers should be careful not to provoke negative reciprocity. Managers are typically extremely concerned to cut wages exactly for this reason. This can explain why wages don't fall during a recession (Bewley 1999). Kube et al. (2008) find a strong negative impact of a wage reduction on workers' productivity in a field experiment. More generally, there is ample evidence that individuals reject unfair offers in bargaining games, albeit at a personal cost (Camerer and Thaler 1995).

Reciprocity, however, does not solely depend on the actions taken and the resulting outcomes, as suggested by the literature cited above. What an individual or organisation intends with a certain action also matters, as evidenced by several experiments, e.g. Brandts and Charness (2003), Falk et al. (2003 and 2008), and Charness and Levine (2007). Some economists have therefore developed models that explicitly take an individual's intentions into account. In the words of Falk and Fischbacher (2006), "a reciprocal action is modeled as the behavioral response to an action that is perceived as either kind or unkind" (pp 294). The perception of the kindness of an action depends on the distributional outcome as well as the intention of the action, see e.g. Rabin (1993), Dufwenberg and Kirchsteiger (2004), and Falk and Fischbacher (2006). These so-called psychological games are highly complex, because the assesment of a player's intentions requires that beliefs are formed over the beliefs of the other player. A more tractable approach is proposed by Levine (1998). In Levine's model, individuals "do not care about whether opponents play "fairly", but rather whether their opponents are nice people" (p. 10) In other words, Levine assumes that individuals are conditionally altruistic: they care more about the welfare of a person when they sense that the person cares for them. This standard game theoretic approach to reciprocity is well in line with the findings reported in organizational psychology and management: when employees infer that their manager or the organization cares about their well-being, they reciprocate with increased commitment, loyalty, and performance (Rhoades and Eisenberger 2002 and Cropanzano and Mitchell 2005).

1.3 Contribution to the literature

The first part of this thesis contributes to the literature on reciprocity in the workplace in various ways. In chapter 2, I use a Levine-type of model to investigate the interaction between reciprocity and pay-for-performance. Although reciprocity can be regarded as a viable substitute for explicit monetary incentives, there is nothing that prohibits employers to provide strong incentives and generous wages at the same time. A few laboratory experiments touch on this issue, showing that explicit incentives are hard to reconcile with reciprocity-induced voluntary cooperation, see e.g. Falk and Kosfeld (2006) and Bartling et al. (2011). However, in these experiments the incentive is either a minimum effort requirement or a fine in case of verified shirking. It is therefore unclear to what extent these results translate to a real-life working situation where workers are risk-averse and employers offer a performance contingent bonus in addition to a base salary.

The theoretical literature on reciprocity in the workplace makes no clear-cut predictions. Englmaier and Leider (2008) find that monetary incentives and a high base salary are substitutes, while Bellemare and Shearer (2011) find that employers who want to give a wage-gift optimally do so by increasing piece rates rather than the base wage. Chapter 2 of this thesis presents a model that contributes to this small literature in various ways. Compared to the existing literature, a key innovation is that my model takes the importance of intentions into account, as described above, and at the same time allows for risk-aversion. Moreover, in chapter 2 I investigate the consequences of worker heterogeneity. A typical finding in laboratory experiments is that not all individuals are equally motivated by reciprocity. For example, in a public-goods game typically 1/3 of experimental subjects does not contribute, regardless of the contribution of others (Fischbacher et al. 2001 and Fischbacher and Gächter 2010). Chapter 2 therefore examines how employers can design a remuneration scheme that not only induces reciprocity, but also screens workers that are reciprocal.

The economic literature so far has understood the employer's treatment of employees mainly in monetary terms. However, employees clearly do not confine reciprocity to the financial aspects of their job. Consider a boss who pays an above-market wage, but hardly pays attention to his employees, and when he does, it is only to insult them because in his opinion a mistake was made. It is hard to imagine that employees are willing to sacrifice their own well-being from time to time in order to favor their generous boss: the opposite seems more likely. Studies in organizational psychology therefore acknowledge the importance of non-tangible, socio-emotional resources for employee motivation (see e.g. Cropanzano and Mitchell 2005). Inspired by this literature, in chapter 3 we assume that workers are reciprocal towards management attention instead of money. We investigate the implications for optimal incentive contracts and confront the theoretical predictions with data from the German Socio-Economic Panel.

The second part of this thesis differs from the first part in two ways. First, we return to the more traditional study of financial incentives as a motivational tool. Second, we use field experiments as primary method of investigation instead of mathematical modelling. Although the experimental method in economics is often used to figure out how social preferences and reciprocity affect human's behavior, it is also a highly valuable tool in the traditional economic framework. The field experiments on the effect of financial incentives on productivity cited earlier are a perfect illustration. Estimations of this effect with naturally occurring cross-section data would be biased, because firms that offer steep financial incentives may do so for a reason we cannot observe, and this reason may be correlated with the firm's productivity. By creating exogenous variation in the data, experiments greatly facilitate establishing causality and hence allow empirical investigation into issues that have received little attention before.

In chapter 4 of this thesis, we address such a long-standing issue, namely how feedback on intermediate relative performance influences subsequent performance. The number of previous studies that address this issue in an economic context is limited to one, namely Casas-Arce and Martinez-Jerez (2009).* Our study adds to the scarce available evidence and uses a unique approach in evaluating the effects of feedback, as will be discussed in more detail below.

In chapter 5 I present a field experiment designed to test some other predictions derived from tournament theory. First, we examine how the distribution of total prize money over the two rounds of an elimination tournament influences performance over the course of the tournament. Second, we investigate whether the incentive effect of the tournament depends on noise in the performance measure. Our study is the first to test these hypotheses within an existing organisation using a natural field experiment. Elimination tournaments have received little attention in the literature, especially in a business context. The effect of prize spread on effort provision is also studied by Altmann et al. (2008), Audas et al. (2004), and DeVaro (2006). Altmann et al. (2008) use a laboratory experiment, while Audas et al. (2004) and DeVaro (2006) both use naturally occurring field data. DeVaro (2006) also examines the effect of noise on the incentive effect of the tournament. Bull (1987) touches on the same issue using a lab experiment. The advantage of a field experiment over naturally occurring data is that there are no confounding factors, hence facilitating reliable causal inference. As compared to a laboratory experiment, the advantage of using a natural field experiment is that it represents a real-life working situation.

1.4 Overview of the thesis

The first part of this thesis analyzes agency problems from various angles. The common theme is that we assume that workers are reciprocal or conditionally altruistic, but the models presented differ considerably from each other.

Chapter 2 of this thesis studies how the base salary and a performance-contingent bonus interact when some workers are conditionally altruistic. Inspired by Levine (1998), I assume that the employer is either altruistic or egoistic, and that a certain fraction of workers are altruistic when they sense that their employer cares for

*Around the time we completed our study, a second paper on relative feedback became available, namely Frank and Obloj (2011).

them. That is, some, but not all, workers are conditionally altruistic. The employer hires one worker from a large pool of workers. He offers a contract consisting of a base salary and a share of the production value. Workers infer the employer's kindness from his contract offer, and those who are reciprocal towards the employer's kindness reciprocate by exerting more effort. I investigate how altruistic employers can optimally distinguish themselves from egoistic employers. I find that altruistic employers offer weak incentives and pay a relatively high base salary. Moreover, I explore how contracts can be designed to signal the employer's kindness and at the same time induce self-selection of workers that are reciprocal towards the employer's kindness. I find that both goals can be attained by offering strong incentives and paying a relatively high base salary.

Chapter 3 differs from chapter 2 and the main body of economic literature by assuming that workers are reciprocal toward management attention instead of money. We analyze a principal-agent model where workers exert effort and managers, in addition to incentive pay, use non-monetary tools of management, which we refer to as management attention. We assume that management attention raises worker's well-being, and is reciprocated with higher effort. We analyze the manager's incentives to give attention and how this affects the optimal provision of incentives for workers. The key finding is that individual pay-for-performance dilutes incentives for the manager to give attention, as paying workers according to their performance reduces the amount at stake for the manager. In a multi-agent setting, this problem can be resolved by offering promotion incentives. We would thus expect that more reciprocal workers are more likely to receive promotion incentives rather than an individual performance bonus. We empirically examine this theoretical prediction using the German Socio-Economic Panel, which contains data on compensation schemes and a measure of an individual's reciprocity. As expected, we find that more reciprocal workers are significantly more likely to receive promotion incentives. However, we do not find a clear negative relation between a worker's reciprocity and the probability of receiving individual pay-for-performance.

The second part of this thesis is devoted to an empirical inquiry into a more traditional economic topic, using a somewhat less traditional method. I present two natural field experiments that test some specific elements of tournament theory within an existing organisation. Chapter 4 examines how intermediate performance feedback affects subsequent performance. Tournament participants who learn during the tournament that they have little chance of winning may give up, while contestants comfortably ahead may slack off, as they expect to win anyway. The relevance of this question is not restricted to tournaments, but can easily be generalized to other non-linear pay-for-performance plans. For instance, consider a salesman who can earn a bonus by attaining a monthly salestarget while receiving weekly or daily sales figures. The bonus cannot be expected to motivate the salesman when intermediate sales are such that the target can be attained easily or when the target is virtually out of reach. Testing this prediction is hard using naturally occurring data due to the lack of a suitable control group: firms typically do not limit the introduction of an incentive scheme to a randomly selected subset of the complete salesforce. A field experiment circumvents this problem, but does not guarantee an unbiased estimate of the presence and strength of dynamic incentive effects. The reason is that performance in a sales context is serially correlated, hence biasing the estimates of the effect of intermediate relative performance on subsequent performance. These complexities have limited the number of previous studies on dynamic incentive effects in an economic context, as noted above.

We contribute to this small literature by taking a unique approach in tackling the issues mentioned above. We conduct a field experiment in a Dutch retail chain, where we introduce a relative performance pay scheme in a randomly selected subset of its 189 stores. Specifically, employees in the treatment stores could win a bonus by outperforming three comparable stores from the control group over the course of four weeks. We provided the treatment stores with weekly feedback on their performance, while the control stores were kept unaware of their involvement. As a result, the performance of the control stores generates exogenous variation in relative performance. By using the performance of the control stores as an instrument, we provide unbiased estimates of the effect of intermediate relative performance on subsequent performance. We find that one percentage point increase in a store's relative performance increases subsequent performance by 0,72%.

The field experiment presented in chapter 5 focuses on elimination tournaments. We run a two-round elimination tournament among a randomly chosen subset of the company's stores. In both rounds of the tournament, participating stores are assigned to poules of four stores that are comparable in terms of historical performance. Half of the stores drop out after the first round, the other half wins a prize and continues to the second round. In the second round, the two best-performing stores again win a prize.

Our primary aim is to investigate the relation between performance and the distribution of total prize money over the two rounds of an elimination tournament. Standard tournament theory predicts that a convex prize structure (i.e. a higher prize in the second round than in the first round, keeping total prize money constant) leads to better second-round performance at the expense of first-round performance. To test this prediction, we run two different treatments that differ by the prize spread only. In the low-spread treatment, prizes are identical in the two rounds, whereas in the high-spread treatment the second-round prize is four times as large as the first-round prize. We find that second-round performance is around 1 percentage point higher in the high-spread treatment as compared to the low-spread treatment, while first-round performance is 0.8 percentage point lower. However, these differences are not statistically significant.

A second objective is to investigate whether the incentive effect of the tournament depends on noise in the performance measure. Theory predicts that noise dilutes incentives to perform, as noise reduces the probability that additional effort leads to a win. To test this prediction, we divide the treatment stores in a high-noise and low-noise group based on the historical variance in performance. We keep both noise groups separated throughout the tournament. Thus, all stores in a given poule are exposed to similar volatility in performance. Our findings are in line with theory: noise negatively influences the positive effect of the tournament on performance. This effect is substantial: an estimated 30% of the stores do not react to the tournament due to high noise.

Chapter 2

Gift-Exchange, Incentives, and Heterogeneous Workers

2.1 Introduction

Economists generally recognise that human's pecuniary motives are not the only determinant of economic outcomes: other considerations play a significant role as well. One of these considerations that has received lots of attention recently is reciprocity, meaning that people are willing to promote the welfare of a kind person and reduce the welfare of an unkind person, even if it comes at a personal cost. The importance of reciprocal motivations in the workplace has been brought under the attention of economists by Akerlof's (1982) seminal paper on the gift-exchange hypothesis. He describes labor contracts as a gift-exchange between employers and employees, where employee's effort and employer's generous treatment of workers are reciprocal gifts.

Generous treatment of employees encompasses several aspects. Of all the aspects mentioned by Akerlof, the wage level has without doubt attracted most attention. The idea that workers reciprocate high wages by exerting more effort has been the subject of considerable empirical examination. Numerous laboratory experiments find a positive relation between employee's effort and the salary offered by the employer, although recent field studies are somewhat less supportive.[1]

[1] An early experimental study is Fehr, Kirchsteiger and Riedl (1993). Fehr and Gächter (2000)

In the light of these substantial efforts to empirically test the gift-exchange hypothesis, the limited amount of theoretical investigation is surprising. For instance, little is known about how employers optimally induce reciprocity when they do not only decide on the wage level, but also on the level of performance pay. Since monetary incentives are an essential element of labour contracts, the perceived generosity of a contract offer may well depend on its incentive intensity.

This chapter studies this question by allowing for reciprocity in an otherwise standard principal-agent model. Thus, the principal decides on a base salary and a piece rate to compensate the worker. Workers are risk-averse and not protected by limited-liability. Although risk-aversion is a common assumption in agency models, its interaction with reciprocity is still an open question. Risk aversion is likely to affect the optimal composition of a gift, as workers are not neutral with respect to the variance of their expected income.

The modeling of reciprocity is inspired by Levine (1998)'s game-theoretic approach. This approach is based on the idea that to many workers, it matters whether their boss cares about them as a person or views them merely as a means to an end. The key assumption is that workers are conditionally altruistic: they care more about the principal's welfare when they sense that the principal cares for them. Specifically, I assume that the principal is either altruistic or selfish, and that conditionally altruistic workers are altruistic to the extent that they believe that the principal is an altruist. Thus, conditionally altruistic workers reciprocate a favor not because of the favor itself, but rather because it signals that the employer cares for them.

An advantage of this approach is that it offers a tractable way to distinguish between authentic and strategic kindness. When workers reciprocate a high wage by exerting more effort, an egoistic principal may want to pay a generous wage for strategic reasons: not because he cares about the workers' well-being, but because he wants to maximize his own profits. This concealed egoism puts the principal's generosity in a completely different light. Workers may therefore not automatically reciprocate high wages, but only when they sense that this kindness is genuine, rather

survey the voluminous literature. See Dur et al. (2008) for an overview of some recent field experiments.

than strategic. There is considerable experimental evidence that people do not only care about distributional outcomes, but are also concerned about the intentions behind an action (see e.g. Charness and Levine, 2007, and the references therein). Apart from the abundant experimental evidence, the relevance of intentions is also apparent in our criminal law system: sentence length depends not exclusively on the harm inflicted on the other party, but to a large extent on whether the harm was caused intentionally or by accident.

Although reciprocity is often considered to have a strong influence on labour contracts (see e.g. Bewley, 1999), a typical finding in laboratory experiments is that not all individuals are equally motivated by reciprocal tendencies. For example, in a three-person gift-exchange experiment by Gächter and Thöni (2010) about 25% of experimental subjects classify as 'egoistic', meaning that they are unwilling to exert effort regardless of the wage they and their colleague receive. Similar patterns are reported by Fischbacher et al. (2001) and Fischbacher and Gächter (2010) in the context of a public goods experiment. This heterogeneity in reciprocity raises the issue whether employers can possibly screen workers, and whether they find it profitable to do so. To investigate this issue, I assume throughout that not all workers are conditionally altruistic: some workers are exclusively motivated by their own material interest. I examine whether an altruistic principal can write a contract that signals his altruism, and at the same time screens conditionally altruistic workers.

The main focus of this chapter is to characterize contracts that induce reciprocity, assuming the principal refrains from screening workers. The first result is that an altruistic principal offers weaker incentives than an egoistic principal, while at the same time increasing the base salary to ensure that he is viewed as an altruist. The reason is that a tight link between pay and performance is not necessary when workers are convinced of the principal's kindness: strong incentives add little to the worker's productivity, but they do expose workers to (unnecessary) risk. Offering strong incentives is therefore suboptimal for a principal who really cares about the worker's well-being. Thus, weak incentives can be considered as part of the altruistic principal's gift to the worker.

The assumption of risk aversion is essential to this result. When workers would be risk-neutral, the principal can resolve all agency problems simply by equating the piece rate with the marginal product. By contrast, worker's risk aversion induces the principal to set the piece rate below the marginal product, and reciprocity can be helpful to further align the interest of the worker and the principal.

The second finding is related to the first and qualifies the standard result, namely that employers induce reciprocity by paying a relatively high expected total compensation. Because part of the altruistic principal's gift is a reduction in incentives, he need not necessarily pay a higher expected total compensation than an egoistic principal to signal his altruism. The reason is that, since an altruistic principal provides workers with little explicit incentives to exert effort, workers' effort may be relatively low, despite their altruistic feelings towards the principal. As a result, pretending to be an altruistic principal is not particularly profitable for an egoistic principal, implying that a relatively low total compensation suffices to distinguish both types. Thus, the interaction between incentives, risk-aversion, and considerations of strategic kindness may divert the usual positive relation between wages and effort. An altruistic principal pays a higher expected total compensation than an egoistic principal only if increased worker motivation leads to sizeable productivity gains, for instance because incentivizing workers via financial incentives is costly due to strong risk aversion.

The third result is that an altruistic principal may find it optimal to write a contract that signals his benevolent intentions, and simultaneously selects conditionally altruistic workers. Perhaps surprisingly, this is accomplished by setting excessively strong incentives and paying a relatively high expected total compensation. The reason for setting strong incentives is that conditionally altruistic workers exert more effort than egoistic workers, and hence have more to gain from an increase in pay-for-performance than egoistic workers. Thus, the paradox is that strong incentives are offered in order to attract the employees who need them the least.

These findings relate to a broad literature on the signaling value of incentives. In particular, the finding that in the absence of screening motives, an altruistic principal induces reciprocity by offering weaker incentives than an egoistic principal, is well in line with recent experimental evidence. Several experimental studies find that incentives are hard to reconcile with reciprocity-induced voluntary cooperation. For instance, a laboratory experiment by Falk and Kosfeld (2006) shows that the implementation of a minimum effort requirement reduces effort, because a considerable group of individuals interpret such an action as a sign of distrust. The implementation of a fine in case effort does not meet a prescribed level has a similar effect, as shown by Fehr and Gächter (2002) and Fehr and List (2004). Their experimental evidence shows that, in the words of Fehr and List (2004, p. 743), "incentives based on explicit threats to penalize shirking backfire by inducing less trustworthy behavior". Ellingsen and Johannesson (2008) provide theoretical underpinnings for this behavior, arguing that control systems and incentives signal that the principal is not worth impressing. Likewise, Sliwka (2007) argues that incentives signal that selfish behavior is the social norm, which demotivates the conformistic agents in the population. These models crucially differ from the model presented here in that esteem and conformism drive the results, instead of reciprocity.

These findings extend to a more natural setting where the principal not only decides on controlling or trusting the agent, but also on a wage level. Falk and Kosfeld (2006) and Bartling et al. (2011) find that when employers have the opportunity to control the agent by limiting his effort discretion, employers either implement a control strategy, which consists of low effort discretion and low wages, or a trust strategy, which consists of high effort discretion and substantial wages. This suggests that paying a high wage and at the same time limiting a worker's effort discretion are conflicting signals. Controlling the worker therefore reduces the effectiveness of a wage gift.[2]

As noted above, this difficulty to reconcile incentives with voluntary cooperation is in line with my model. An important difference, however, is that I look at another

[2] This phenomenon is also referred to as partial-crowding out (Fehr and Gächter, 2002): keeping the wage constant, voluntary cooperation is lower when (stronger) incentives are implemented, where voluntary cooperation is defined as the difference between actual and privately optimal effort. Fehr and Gächter (2002) and Bellemare and Shearer (2011) provide evidence in line with the partial-crowding hypothesis.

In my model, the reason for crowding is that a high piece rate diminishes the share of the marginal product that accrues to the principal, and therefore restricts the worker's opportunities to reciprocate the principal's favourable treatment. Hence, a tight link between effort and monetary reward reduces the principal's return to signaling altruism.

type of incentives, namely piece rates instead of minimum effort requirements. As a result, in my model incentives are not necessarily viewed negatively. This allows altruistic employers to signal their altruism and at the same time screen workers by offering relatively strong incentives and a high expected total compensation.

This chapter proceeds as follows. The next section discusses the related literature. Section 2.3 sets out the model. Section 2.4 analyzes the observable types case, which serves as a benchmark for the analysis in section 2.5 where types are assumed to be unobservable. There, I first show that pooling equilibria do not exist, and then proceed to the analysis of separating equilibria. Finally, in section 2.6 I conclude and provide some avenues for further research.

2.2 Related literature

In economics, several authors have suggested ways of modeling reciprocity and underlying intentions, for example Rabin (1993), Levine (1998), Dufwenberg and Kirchsteiger (2004), Falk and Fischbacher (2006) and Battigalli and Dufwenberg (2009). Our model is based on Levine's approach because it provides a tractable and natural way to model the findings reported in organizational psychology and management: when employees infer that their manager or the organization cares about their well-being, they reciprocate with increased commitment, loyalty, and performance (see, for example, the reviews by Rhoades and Eisenberger 2002 and Cropanzano and Mitchell 2005). As pointed out by Gul and Pesendorfer (2010), this approach facilitates analysis and allows to distinguish between genuine kindness and instrumental kindness.

Despite the recent attention for modeling reciprocal behavior in the workplace, there are only few studies that investigate the relation between reciprocity and monetary incentives: Englmaier and Leider (2008), Arbak and Kranich (2005), and Bellemare and Shearer (2011). Compared to these studies, the key innovation is that I account for worker's risk-aversion as well as a concern for the intentions of the employer. The interaction between these assumptions plays a critical role when deriving the employer's optimal choice of base salary and bonus payments.

Englmaier and Leider (2008) incorporate reciprocity in a principal-agent model with risk averse agents. Their main finding is that incentive pay and reciprocal motivations are substitutes, which is qualitatively in line with my results. The key difference is that Englmaier and Leider largely ignore the importance of intentions, and instead assume that positive reciprocity is automatically induced when agents expect to receive a rent.

A working paper by Arbak and Kranich (2005) is closely related in the sense that incomplete information and Levine-type conditional altruism are key features of their modeling set-up. They assume that employers signal their type by increasing piece rates, but fail to analyze the case where employers can also resort to the base salary as a signaling tool.

In a related model, Bellemare and Shearer (2011) address this shortcoming. They find that employers induce reciprocity by increasing piece rates rather than by increasing the base salary. The main reason is that piece rates have a direct incentivizing effect, in addition to the effect of the reciprocity induced by the gift. This result follows from their assumption that workers are risk neutral and that a limited-liability constraint is always binding. Therefore, in the absence of a signaling motive piece rates are suboptimally low. As higher piece rates bring workers' incentives closer to the socially optimal level, the most efficient way to signal kindness is by increasing incentives. By contrast, in the model presented in this chapter stronger incentives expose workers to more risk, which is costly. Thus, assuming that workers are risk averse and unconstrained by limited-liability crucially affects the optimal composition of a wage gift.

Chapter 3 of this thesis also studies the relation between reciprocal motivations and incentive pay. However, the model presented in chapter 3 does not allow for monetary gift-exchange, but focuses on social gift-exchange instead, meaning that the resources of the gift-exchange are non-monetary.[3] In particular, in chapter 3,

[3] As argued by Akerlof (1982), manager's benevolent treatment of workers encompasses several aspects. Studies in management and organizational psychology distinguish between two broad categories: economic resources and socioemotional resources. Economic resources address "financial needs and tend to be tangible", whereas socioemotional resources address "social and esteem needs (and are often symbolic and particularistic)." (Cropanzano and Mitchell 2005, p. 881) In economics, monetary gift-exchange refers to the former, whereas social gift-exchange in the chapter 3 refers to the latter.

the manager's gift exclusively consists of management attention. Another limitation is that the model in chapter 3 does not allow for strategic kindness of the principal. Both limitations are addressed by Dur (2009), but his model does not allow for incentive pay.

Another topic I address is how contracts can select reciprocal worker types and at the same time signal the principal's altruism. Interestingly, Kosfeld and Von Siemens (2009, 2011) investigate how employers use a combination of base salary and performance pay to screen reciprocal workers. An important difference is that they concentrate on reciprocity between colleagues who work in a team, rather than between worker and principal. The principal's contract offer therefore does not signal private information.

Sliwka (2007) does not study reciprocity, but he also studies contracts that signal the principal's private information and at the same time screen workers. As noted above, the main idea in his paper is that incentives signal that selfish behavior is the social norm, which demotivates the conformistic agents in the population. In addition to this signaling effect, incentives can also screen worker types when selfish and altruistic workers differ in their preferences over incentive intensity. Hence, the optimal decision whether to trust or to incentivize agents takes both the signaling and selection effect into account. An important difference with my analysis, however, is that in Sliwka the principal faces the binary decision whether to trust or incentivize workers, whereas in my model the principal has two continuous instruments (wage and base salary) at his disposal. Therefore, the principal has richer opportunities to simultaneously signal to and screen workers.[4]

2.3 The model

I consider a risk-neutral principal who is hiring one worker from a large population of workers. Workers are risk-averse and some of them are conditionally altruistic, meaning that the extent to which the worker is altruistic depends on the principal's

[4] To the best of my knowledge, the only other paper that designs a contract that signals information to and at the same time screens the other contracting party, using two continuous instruments, is Soberman (2003).

altruism. The expected utility of a worker of type i is described by:

$$u_i = -\exp^{-r\left[b(e+\varepsilon)+s-\frac{1}{2}\theta e^2+\gamma_i(\alpha_j)E(\pi_j)\right]}. \tag{2.1}$$

This specification is widely used to describe the utility of risk-averse agents, where $r > 0$ captures the extent of risk-aversion. Production is simply given by effort e, but is prone to random shocks ε that are normally distributed with variance σ^2. Effort is non-contractible, but assuming that output can be observed, the worker earns a share b of observed output $(e+\varepsilon)$ and a base salary denoted by s. The worker's costs of effort are represented by $\frac{1}{2}\theta e^2$. Furthermore, $\gamma_i(\alpha_j)$ captures the strength of the worker's altruism towards his principal, which positively depends on the principal's altruism α_j. Note that the worker cares about the principal's expected payoff $E(\pi_j)$ instead of his actual payoff π_j. The reason is that it would be nonsensical to assume that the worker is risk-averse over the payoff of a risk-neutral principal. Workers differ in their altruism function $\gamma_i(\alpha_j)$, where $0 \leq \gamma_i(\alpha_j) < 1$. I distinguish between a worker's 'type' and a worker's 'altruism'. A worker's 'type' refers to his altruism function, whereas his 'altruism' refers to the outcome of the function $\gamma_i(\alpha_j)$. The typespace is specified below.

As is common in the literature, I remove the uncertainty on ε from the worker's utility function by deriving the certainty equivalent, which allows for convenient transformation of the utility function into:[5]

$$E(u_i) = be + s - \frac{1}{2}\theta e^2 - \frac{1}{2}r\sigma^2 b^2 + \gamma_i(\alpha_j)E(\pi_j). \tag{2.2}$$

An increase in the bonus exposes the worker to more risk, which subtracts from the positive effect of the bonus on the worker's utility. To ensure that an increase in the bonus does not expose the worker to so much additional risk that his utility is

[5] This standard transformation is only correct when there is no uncertainty on the principal's type, either because types are observable or because the worker puts all probability mass on a certain type. Uncertainty on the principal's type reduces the worker's welfare compared to the utility suggested by (2.2). Because, as we shall see, only separating equilibria exist, there is no uncertainty on the principal's type in equilibrium, and hence it is safe to ignore the worker's 'preference for certainty' in this domain for simplicity. Taking this preference into account would only strengthen the results, because deviation from the equilibrium strategy is less attractive when it leads to uncertainty about the principal's type.

lowered, I assume throughout that $r\sigma^2\theta \leq 1$.[6]

The expected payoff of a principal of type j is described by

$$E(\pi_j) = (1-b)e - s + \alpha_j E(u_i), \tag{2.3}$$

where $\alpha_j E(u_i)$ captures the altruistic feelings of the principal.[7] Analogous to the workers, I assume that $0 \leq \alpha_j < 1$. Principals differ in type α_j, implying that the worker's altruism depends on the specific match between the worker's and the principal's type.

Types are private information. There are two types of workers and principals, namely selfish (l) and (conditionally) altruistic (h) types. A selfish or egoistic principal has no altruistic feelings at all: $\alpha_l = 0$. An altruistic principal's care for the worker is denoted $\alpha_h \in (0,1)$. The prior probability that the principal is selfish is given by μ. Because the principal's type is unobservable, a worker's altruism depends on his beliefs concerning the principal's type. Workers update their beliefs after observing the principal's contract choice using Bayes' rule. When applicable, I rule out unreasonable beliefs by requiring that beliefs satisfy the intuitive criterion.[8]

Workers are either egoistic or conditionally altruistic/reciprocal. Specifically, I assume that an egoistic worker type never takes the principal's welfare into account regardless of the principal's altruism ($\gamma_l(\alpha_j) = 0$ for any α_j) and that a reciprocal worker type is completely egoistic when he believes that the principal is selfish ($\gamma_h(\alpha_l) = 0$). Thus, only a conditionally altruistic worker who believes that he is employed by an altruistic principal takes the welfare of the principal into account,

[6] This assumption does not affect the main results. The mathematical condition is derived in the proof of lemma 4 given in the appendix.

[7] Obviously, the coefficients α_j and $\gamma_i(\alpha_j)$ depend on the units in which utility is measured, because utility must be measured in interpersonally comparable units. Therefore, the principal's payoff function only makes sense when the worker's utility is measured by the transformed utility function (2.2). Note that both players take the other's total payoff into account, so including the immaterialistic part of the payoff function. This assumption is not essential for the qualitative results and is not uncommon in the literature, see e.g. Becker (1974), Barro (1974), Bernheim and Stark (1988), Ley (1997), Sol (2010) and Dur and Sol (2010). Confining altruism to the material payoffs, as in Levine (1998), leads to a discrepancy between the worker's actual utility and what the principal believes the worker's utility to be. This is inconvenient analytically.

[8] Beliefs satisfy the intuitive criterion if, for all out-of-equilibrium actions, zero probability is assigned to player types that can only lose compared to their equilibrium payoff, see Cho and Kreps (1987).

i.e. $\gamma_h(\alpha_h) > 0$. I assume that a fraction δ of all workers in the population is selfish, whereas the remaining fraction $(1 - \delta)$ is conditionally altruistic.

The timing of the game is as follows. First, the principal decides on a remuneration scheme (b, s). A worker accepts the contract if it yields him an expected utility of at least his reservation utility $\overline{u}$. Hence, as is common in principal-agent models, workers have no bargaining power. Also, I assume that $\overline{u}$ does not depend on $\gamma_i(\alpha_j)$: the value of a worker's outside option does not depend on his type. In case none of the workers accepts the contract, the principal's payoff is zero. Finally, the employed worker decides on his effort level e.

2.4 Analysis when types are observable

In this section, I assume that both players learn about each other's type before they make any decision. The reason for studying this case is that it yields some insights that will be valuable later on. I solve for a subgame perfect equilibrium using backward induction. The worker's effort choice follows from maximization of his utility function (2.2). The first-order condition is described by:

$$b - \theta e + \gamma_i(\alpha_j)(1 - b) = 0. \tag{2.4}$$

It is instructive to see what happens if $b = 0$ or if $b = 1$. If $b = 0$, the worker only exerts effort out of an altruism motive. By contrast, if the worker is residual claimant $(b = 1)$, the worker's actions do not affect the principal's profits. Therefore, any worker type equates the marginal costs of effort with the marginal product $(\theta e = 1)$, independent of his altruistic feelings. Rewriting the first-order condition (2.4) gives the worker's optimal effort choice e_i^*:

$$e_i^* = \frac{b + \gamma_i(\alpha_j)(1 - b)}{\theta}. \tag{2.5}$$

It can easily be seen that effort increases in financial incentives for any b and in the worker's altruism as long as $b < 1$. Also, it is easily verified that altruism reduces the motivational effect of financial incentives $(\frac{de}{db})$ and vice versa: financial incentives reduce the responsiveness of effort towards altruism ($\frac{de}{d\gamma}$). The latter effect is intuitive: the larger the share of the marginal product that accrues to the worker, the smaller the share that accrues to the principal, hence the smaller the worker's possibilities to increase the principal's welfare. Therefore, the model predicts partial crowding-out of voluntary cooperation.[9] The negative effect of altruism on the motivational effect of financial incentives follows from the fact that the more the worker cares for his boss, the less he enjoys his bonus. In the extreme case that $\gamma_i(\alpha_j)$ approaches 1, the worker cannot be motivated by incentive pay because he cares about the principal's payoff as much as he cares about his own payoff.

The principal's choice of the optimal bonus b follows from maximization of his expected payoff, where he takes into account the worker's response to financial incentives and the worker's participation constraint:

$$\max_{s,b} \; E(\pi_j) = (1-b)e_i^* - s + \alpha_j E(u_i)$$
$$\text{s.t. } E(u_i) = be_i^* + s - \frac{1}{2}\theta e_i^{*2} - \frac{1}{2}r\sigma^2 b^2 + \gamma_i(\alpha_j)E(\pi_j) \geq \overline{u}.$$

Since the principal cares more about his own payoff than about the worker's utility ($0 \leq \alpha_j < 1$), it is not optimal to leave a rent to the worker. The principal thus reduces the base salary until the participation constraint binds. Inserting the base salary implied by the participation constraint into the objective function and differentiating with respect to b, we obtain the following first-order condition for optimal incentive provision:

$$\frac{de_i^*}{db}(1-\theta e_i^*) - r\sigma^2 b = 0.$$

This condition elucidates the principal's trade-off. An increase in the bonus has one benefit and two costs. The benefit is that an increase in the bonus leads to additional effort, which contributes to the principal's payoff with the size of the marginal product. However, because the worker's participation constraint is binding, the worker needs to be compensated for the additional cost of providing effort. Moreover, risk-

[9] Voluntary cooperation is defined as the difference between actual effort (e_i^*) and privately optimal effort ($\frac{b}{\theta}$). Clearly, $\frac{\gamma_i(\alpha_j)(1-b)}{\theta}$ is decreasing in the bonus. Note that similar expressions can be found in Arbak and Kranich (2005) and Sliwka (2007).

averse workers need to be compensated for exposure to income uncertainty. We derive the payoff-maximizing bonus b^* by inserting (2.5) and its derivative into the first-order condition, which yields:

$$b^* = \frac{[1-\gamma_i(\alpha_j)]^2}{[1-\gamma_i(\alpha_j)]^2 + \theta r \sigma^2}. \tag{2.6}$$

Clearly, the bonus decreases in the worker's altruism $\gamma_i(\alpha_j)$. The reason is twofold. First, the more altruistic the worker, the smaller the motivational effect of financial incentives. In terms of the first-order condition, $\frac{de_i^*}{db}$ decreases in $\gamma_i(\alpha_j)$, hence reducing the marginal benefit of the bonus. The second reason is that the more altruistic the worker, the more effort he exerts, and consequently the higher his marginal costs of effort. Because the worker needs to be compensated for his costs of effort, it is more costly to stimulate effort further using financial incentives. Note that when the worker has no altruistic feelings towards his principal ($\gamma_i(\alpha_j) = 0$), we obtain the standard solution: the bonus is decreasing in the worker's risk aversion and in the variance of the error term.

A result that is harder to anticipate is the ambiguous total effect of an increase in $\gamma_i(\alpha_j)$ on effort. Given the level of financial incentives, worker's altruism has a positive effect on effort, but this is possibly more than offset by the corresponding decrease in financial incentives. Specifically, the total derivative $\frac{de_i^*}{d\gamma_i(\alpha_j)}$ is described by

$$\frac{de_i^*}{d\gamma_i(\alpha_j)} = \frac{r\sigma^2\left(\theta r\sigma^2 - [1-\gamma_i(\alpha_j)]^2\right)}{\left([1-\gamma_i(\alpha_j)]^2 + \theta r\sigma^2\right)^2},$$

which is positive if $\theta r\sigma^2 > [1-\gamma_i(\alpha_j)]^2$. Thus, an increase in $\gamma_i(\alpha_j)$ only has a positive effect on effort when workers are relatively risk-averse. Strong risk-aversion implies that the bonus is relatively small ($b^* < \frac{1}{2}$). As a result, additional effort has a relatively large effect on the principal's welfare, and hence the worker's effort is more responsive to changes in $\gamma_i(\alpha_j)$. By contrast, when workers are hardly risk-averse, the bonus is close to the marginal product, which restricts the worker's opportunities to benefit the principal. The negative effect of a reduction in incentives therefore outweighs the positive effect of an increase in worker's altruism.

The base salary is such that the worker's participation constraint is exactly satisfied:

$$s = \overline{u} - b^* e_i^* + \frac{1}{2}\theta e_i^{*2} + \frac{1}{2} r\sigma^2 b^{*2} - \gamma_i(\alpha_j)E(\pi_j). \tag{2.7}$$

As long as the principal's expected payoff is positive[10], total compensation is decreasing in the worker's altruism. The main reason is a compensating wage differential: the more utility workers derive from the non-monetary aspects of their job, the lower the required monetary compensation. Altruism leads to a rent that can be extracted when types are observable. A second reason is that the strength of financial incentives decreases in the worker's altruism, which is reflected in a lower risk-premium when workers are risk averse.

Worker's altruism has a positive effect on profits, for two reasons. First, because preferences are more aligned, altruistic workers are willing to exert more effort, given the strength of financial incentives (see (2.5)). Second, as more altruistic workers derive more utility from the non-monetary aspects of their job, they are willing to accept a lower base salary, keeping b constant (see (2.7)).

2.5 Analysis when types are unobservable

In the previous section we obtained some basic insights by assuming that types are observable. In this section I study the more interesting case where types are private information. Because neither the worker's type nor the principal's type is observable, contracts potentially have a dual role of both signaling the principal's type and screening worker's types. I consider four possible Perfect Bayesian equilibria, namely pooling and separating equilibria, where in each equilibrium the principal either screens or abstains from screening. In pooling equilibria, the equilibrium contract offer is independent of the principal's type. By contrast, in separating equilibria different types offer different contracts.

In the next two subsections, I will show that pooling equilibria do not exist. Then, I analyze a fully separating equilibrium where the principal signals his altru-

[10] Firms can only stay in business as long as expected profits are at least zero. It is easily verified that profits are increasing in the worker's altruism, so an altruistic principal obtains a positive expected payoff even when an egoistic principal earns zero profits in expectation.

ism, but abstains from screening worker types. Finally, I study a fully separating equilibrium where the altruistic principal writes a contract that signals his altruism and simultaneously screens worker types. I abstract from mixed strategy equilibria. Hence, throughout I refer to a fully separating equilibrium simply as separating equilibrium.

2.5.1 Analysis: Pooling equilibrium

This section studies existence of a pooling equilibrium where the principal does not screen workers. I start the analysis by noting that any equilibrium contract, pooling or separating, has to satisfy four constraints. The first two constraints consist of each principal's incentive compatibility constraint (ICC): neither of the types should have an incentive to deviate. The other two constraints are the participation contraints (PC) of both worker types. In equilibria where principals do not screen workers, both worker types should be willing to accept the equilibrium contract(s). By contrast, screening of workers' types requires that the equilibrium contract is acceptable for conditionally altruistic workers only. In equilibrium, workers' beliefs should be correct. Thus, in a pooling equilibrium, workers believe that the principal is egoistic with probability μ.

In order to show that a pooling equilibrium does not exist, it is instructive to compare the profit functions of the egoistic and altruistic principal, see equation (2.3). The crucial difference between the two types is that the altruistic principal enjoys the expected utility of the worker, in addition to his monetary payoff. Whether a deviation from the equilibrium contract is profitable depends to a large extent on reciprocal workers' out-of-equilibrium beliefs. When reciprocal workers interpret deviation as a signal that the principal is egoistic, deviation is not profitable for both types of principals. By contrast, if deviation is interpreted as a signal of altruism, both types are willing to deviate.

The intuitive criterion, however, restricts the set of possible out-of-equilibrium beliefs. In particular, it rules out that the worker believes that a deviation comes from an egoistic principal when that deviation cannot be profitable for an egoist, while an altruist gains by such a deviation. Such a contract always exists. For instance, the altruistic principal can deviate by increasing the base salary with an amount equal to the additional revenues he gains when the worker believes with probability one that the contract is offered by an altruist. The egoist can only lose by offering such a contract. The altruistic principal, however, gains by offering such a contract, because his monetary payoff is unaffected, while at the same time the worker obtains a strictly higher utility, which is valuable for an altruist. Thus, the intuitive criterion dictates that such a deviation must come from an altruistic principal, implying that the altruistic principal always has an incentive to deviate. The following proposition summarizes the discussion above:

Proposition 2.1 *A pooling equilibrium where the principal does not screen workers, does not exist.*

2.5.2 Analysis: Pooling and screening equilibrium

Instead of pooling on a contract that is acceptable for both types of workers, principals may also pool on a contract that is acceptable for conditionally altruistic workers only. As noted above, screening requires that an egoistic worker's participation constraint is violated, while at the same time satisfying the other constraints. The worker's utility function (2.2) reveals that a reciprocal worker derives more utility from a given contract than a selfish worker, as long as he believes with positive probability that the principal is an altruist. Thus, the principal can screen workers by extracting (part of) reciprocal workers' rent. The possibility to screen workers, however, does not change the altruistic principal's incentive to break the proposed equilibrium:

Proposition 2.2 *A pooling equilibrium where the principal screens workers does not exist.*

Proof. See appendix. ■

The intuition is the same as in a pooling equilibrium where the principal does not screen workers: the altruistic principal always finds it profitable to reveal his type. He deviates by offering a contract that cannot be profitably offered by an egoistic principal, and by the intuitive criterion, such a deviation must be interpreted as a credible signal of his altruism. The altruistic principal enjoys the resulting increase in worker's utility, while earning the same (or slightly lower) monetary profits. A formal proof is provided in the appendix, as the proof is complicated by the screening constraint outlined above.

2.5.3 Analysis: Separating equilibria

This section studies separating equilibria: equilibria where the altruistic principal offers a contract that signals his altruism. As pointed out above, in a separating equilibrium the principal either screens or refrains from screening workers. I will distinguish between the two cases after laying out the common structure of these equilibria.

I start the analysis by deriving the egoistic principal's contract choice. Using the assumption that a selfish principal does not care about worker's utility, his expected payoff (see equation (2.3)) can be written as:

$$E(\pi_l) = (1 - b_l)\left[\delta e_{ll} + (1 - \delta)\, e_{hl}\right] - s_l, \tag{2.8}$$

where the subscript (l) is used to indicate that the remuneration scheme (b_l, s_l) is offered by a selfish principal. Similarly, e_{ij} denotes the effort of a worker of type i who faces the incentive scheme offered by a principal of type j and consequently believes that he is employed by a principal of type j.[11] His effort choice is described by equation (2.5). Because beliefs should be correct in equilibrium, a worker of type i observing the contract (b_l, s_l) correctly believes that he is employed by a selfish principal, implying that his expected utility (2.2) is described by:

$$E(u_i) = b_l e_{il} + s_l - \frac{1}{2}\theta e_{il}^2 - \frac{1}{2} r\sigma^2 b_l^2 + \gamma_i(\alpha_l) E(\pi_l) \geq \overline{u}. \tag{2.9}$$

Clearly, because $\gamma_i(\alpha_l) = 0$ irrespective of a worker's type, both worker types exert the same effort and derive the same utility from accepting the selfish principal's

[11] Because in equilibrium beliefs are always based on the observed contract offer, this shorthand notation suffices to describe the equilibrium contracts.

equilibrium contract. This implies that both types require the same compensation to satisfy their participation constraint. Because in any separating equilibrium a selfish principal has no reason to screen or to signal his type, he does not distort his optimal contract choice compared to the case when types are observable. Thus, the following lemma applies:

Lemma 2.1 *In any separating equilibrium, the egoistic principal offers a bonus* $b_l = \frac{1}{1+r\sigma^2\theta}$ *and a base salary that exactly satisfies the worker's participation constraint:*

$$s_l = \overline{u} - b_l e_{il} + \frac{1}{2}\theta e_{il}^2 + \frac{1}{2}r\sigma^2 b_l^2. \tag{2.10}$$

Proof. To show that in any separating equilibrium the egoistic principal always finds it optimal to offer the contract specified above, first note that in any separating equilibrium screening makes no sense for the egoistic principal and is not even possible, as $\gamma_i(\alpha_l) = 0$ (see (2.9)). Thus, all possible equilibrium contracts satisfy the worker's PC as defined by (2.9). To show that the contract specified above is the unique equilibrium contract, consider any arbitrary separating equilibrium. As worker's beliefs are correct in equilibrium ($\gamma_i(\alpha_l) = 0$), the selfish principal cannot be confronted with a reduction in worker's altruism when deviating. Therefore, of all possible equilibrium contracts, the egoistic principal will always choose the contract that maximizes his expected payoff. The properties of this contract follow from the analysis when types are observable and $\gamma_i(\alpha_l) = 0$. ■

Although the egoistic principal's equilibrium contract choice is highly intuitive, it cannot be sustained for all out-of-equilibrium beliefs. For instance, a reduction in the base salary should be viewed as coming from an egoistic principal, otherwise the egoistic principal can gain by slightly reducing the base salary, hence selecting conditionally altruistic workers. It seems natural to assume that a reduction in the base salary does not lead to more positive beliefs. Similarly, a salary raise should not be regarded so optimistic that the raise is actually profitable. Formally:

Lemma 2.2 *In any separating equilibrium, the worker's out-of-equilibrium belief upon observing an alternative contract offer* (b', s')*, is defined as follows. Let* $E(\pi_l^*) = E(\pi_l(b_l, s_l))$ *denote the egoistic principal's expected payoff in any separating equi-*

librium. Let $pr(\alpha_j = \alpha_h)$ denote the worker's belief. There are two cases:
1) When (b', s') is such that $E(\pi_l(b', s') \mid pr(\alpha_j = \alpha_h) = 1) \leq E(\pi_l^*)$, the worker believes that $pr(\alpha_j = \alpha_h) = 1$.
2) When (b', s') is such that $E(\pi_l(b', s') \mid pr(\alpha_j = \alpha_h) = 1) > E(\pi_l^*)$, then $pr(\alpha_j = \alpha_h)$ should be such that $E(\pi_l(b', s') \mid pr(\alpha_j = \alpha_h)) < E(\pi_l^*)$. Such a belief always exists.

Proof. The first condition directly follows from the requirement that beliefs satisfy the Intuitive Criterion. The second condition ensures that the egoistic principal has no profitable deviation. To show that such a belief always exists, notice that $pr(\alpha_j = \alpha_h) = 0$ implies that $E(\pi_l(b', s') \mid pr(\alpha_j = \alpha_h) = 0) < E(\pi_l^*)$. ■

Deriving the altruistic principal's equilibrium contract is more difficult. In a separating equilibrium, any contract that cannot be profitably offered by an egoistic principal is a possible equilibrium contract. However, it is easily verified that the intuitive criterion rules out all equilibria in which the altruistic principal offers a contract that does not maximize his expected payoff.[12] Therefore, of all possible equilibrium contracts, the altruistic principal always chooses the contract that maximizes his expected payoff.

To derive this contract, it is instructive to inspect the altruistic principal's payoff function. Assuming that the equilibrium contract (b_h, s_h) succeeds in credibly signaling the principal's altruism, rewriting equation (2.3) yields:

$$E(\pi_h) = (1 - b_h)\left[\delta e_{lh} + (1 - \delta) e_{hh}\right] - s_h + \alpha_h \left[\delta E(u_l) + (1 - \delta) E(u_h)\right]. \quad (2.11)$$

This equation shows that the altruistic principal's payoff positively depends on a worker's altruism for two reasons. First, as long as $b_h < 1$, a reciprocal worker will put more effort into his job than a selfish worker ($e_{hh} > e_{lh}$). Second, a worker's expected utility $E(u_i)$ increases in his altruism, which is valuable for a principal who has altruistic feelings. For these two reasons, the altruistic principal may benefit from writing a contract that convinces reciprocal worker types that he is an altruist. In addition, he may find it profitable to screen workers, which means that $\delta = 0$. In

[12] The reason is that, since the equilibrium contract by definition satisfies the requirement that it cannot profitably be offered by an egoistic principal, the intuitive criterion dictates that the contract must be offered by an altruistic principal.

order to credibly signal altruism, the equilibrium contract (b_h, s_h) should satisfy two incentive compatibility constraints: the selfish principal should have no incentive to mimic the altruist and vice versa:

$$(1-b_l)e_{\delta l} - s_l \geq (1-b_h)e_{\delta h} - s_h, \tag{ICC1}$$

$$(1-b_l)e_{\delta l} - s_l + \alpha_h \overline{u} \leq (1-b_h)e_{\delta h} - s_h + \alpha_h \left[\delta E(u_l) + (1-\delta)E(u_h)\right], \tag{ICC2}$$

where $e_{\delta l} = \delta e_{ll} + (1-\delta)\, e_{hl}$ and $e_{\delta h} = \delta e_{lh} + (1-\delta)\, e_{hh}$. The incentive compatibility constraints in case the altruistic principal screens workers assume that $\delta = 0$ on the right hand side of the inequalities. It is essential to note that when ICC1 is satisfied, ICC2 can only be satisfied if the difference $\alpha_h \left[\delta E(u_l) + (1-\delta)E(u_h)\right] - \alpha_h \overline{u}$ is large enough. This observation reveals why in equilibrium a principal with altruistic feelings is willing to engage in costly signaling: not because he earns a higher monetary payoff, but because he values workers' utility. The altruistic principal's monetary payoff is constrained by the egoistic principal's monetary payoff: an egoistic principal will imitate any contract that yields a higher monetary payoff than his own equilibrium contract. Therefore, any positive difference in effort $(e_{\delta h} - e_{\delta l})$ must be reflected in higher payments to the worker. Reasoning further along these lines, the following lemma does not come as a surprise:

Lemma 2.3 *The incentive compatibility constraint ICC1 is binding in any separating equilibrium, implying that ICC2 is slack.*

Proof. Suppose, per absurdum, that ICC1 is slack. The altruistic principal chooses the contract that maximizes his payoff. Depending on whether he screens or not, the contract satisfies the participation constraint of both workers or that of conditionally altruistic workers only. Since this contract maximizes his payoff, ICC2 is satisfied, but ICC1 is violated. To see this, recall that profits are increasing in the worker's altruism, because more altruistic workers exert more effort (see (2.5)) and are willing to accept a lower base salary (see (2.7)). Therefore, the egoistic principal finds it profitable to mimick the altruistic principal. Hence, ICC1 binds in any separating equilibrium. Because any equilibrium contract satisfies $E(u_i) \geq \overline{u}$ for the worker who accepts the contract, ICC2 is slack when ICC1 holds with equality. ■

Since ICC1 is binding, in any separating equilibrium the optimal contract will be a (b_h, s_h) −combination such that ICC1 holds with equality and that ICC2 is slack.[13] I will now proceed to discuss the separating equilibrium where the principal does not screen.

Separating equilibrium: no screening

When the altruistic principal does not screen workers, his equilibrium contract should satisfy the participation constraint (PC) of both worker types. For ease of reference, the selfish worker's PC is described by

$$b_h e_{lh} + s_h - \frac{1}{2}\theta e_{lh}^2 - \frac{1}{2} r\sigma^2 b_h^2 \geq \overline{u}. \tag{PCL}$$

Assuming that the contract (b_h, s_h) credibly signals the principal's altruism, a reciprocal worker's PC is described by

$$b_h e_{hh} + s_h - \frac{1}{2}\theta e_{hh}^2 - \frac{1}{2} r\sigma^2 b_h^2 + \gamma_h(\alpha_h) E(\pi_h) \geq \overline{u}. \tag{PCH}$$

Comparison of these two constraints reveals that a reciprocal worker derives more utility from a given equilibrium contract (b_h, s_h) than a selfish worker. This implies that when a selfish worker's PC is satisfied, a reciprocal worker's PC is also satisfied. Because the equilibrium contract (b_h, s_h) by assumption satisfies the selfish worker's PC, the reciprocal worker's PC cannot be binding in the proposed equilibrium.

The problem is to find the contract that maximizes the altruistic principal's expected payoff, provided the four constraints outlined above are satisfied. Because PCH and ICC2 are both slack, the problem can be reformulated in a convenient way. As the altruistic principal's monetary payoff is constrained by the profits earned by the egoistic principal, he maximizes his total payoff by choosing a contract that exactly satisfies ICC1 and maximizes the expected utility of the worker. Denoting the egoistic principal's equilibrium profits as $E(\pi_l^*)$, the maximization problem can

[13] Obviously, the altruistic principal should offer a different contract than the egoist, i.e. the contract $(b_h = b_l, s_h = s_l)$ is not feasible. As we shall see, this condition is always satisfied if $r\sigma^2 > 0$.

thus be formulated as:

$$\max_{s_h, b_h} \quad E(\pi_h) = E(\pi_l^*) + \delta E\left(u_l\right) + (1-\delta)\, E(u_h)$$

$$\text{s.t.} \quad s_h = (1-b_h)e_{\delta h} - E(\pi_l^*), \tag{ICC1}$$

$$\text{s.t.} \quad b_h e_{lh} + s_h - \frac{1}{2}\theta e_{lh}^2 - \frac{1}{2} r\sigma^2 b_h^2 \geq \overline{u}. \tag{PCL}$$

Intuitively, a reasonable conjecture is that maximization of worker's utility ensures that the selfish worker's PC is satisfied. For ease of exposition, I assume that PCL is satisfied and show afterwards that this conjecture is correct. This allows me to rewrite the problem by substituting $E\left(u_l\right)$, $E(u_h)$, and s_h, yielding the following:

$$\max_{b_h} \quad E(\pi_h) = \delta \left(e_{lh} - \frac{1}{2}\theta e_{lh}^2 - \frac{1}{2} r\sigma^2 b_h^2 \right) +$$
$$(1-\delta) \left(e_{hh} - \frac{1}{2}\theta e_{hh}^2 - \frac{1}{2} r\sigma^2 b_h^2 + \gamma_h(\alpha_h) E(\pi_h) \right).$$

Taking the derivative to b_h gives an insightful first-order condition:

$$\delta \left(\frac{de_{lh}}{db_h} \left(1-\theta e_{lh}\right) - r\sigma^2 b_h \right) + (1-\delta) \left(\frac{de_{hh}}{db_h} \left(1-\theta e_{hh}\right) - r\sigma^2 b_h \right) = 0.$$

The first-order condition is the same as in the observable types case, but weighted according to the prevalence of the two worker types. An altruistic principal thus offers the bonus that maximizes the expected surplus. Therefore, depending on the fraction of selfish workers in the population (δ), the payoff-maximizing bonus b_h lies between $b^* = \frac{[1-\gamma_h(\alpha_h)]^2}{[1-\gamma_h(\alpha_h)]^2 + \theta r\sigma^2}$ and $b_l = \frac{1}{1+\theta r\sigma^2}$, implying that an altruistic principal offers a lower bonus than a selfish principal. This can also be seen after rewriting the first-order condition:

$$b_h = \frac{\delta + (1-\delta)\left[1-\gamma_h(\alpha_h)\right]^2}{\delta + (1-\delta)\left[1-\gamma_h(\alpha_h)\right]^2 + \theta r\sigma^2}.$$

The finding that $b_h < b_l$ implies that an altruistic principal pays a higher base salary than an egoistic principal. Recall the assumption made earlier that $r\sigma^2\theta \leq 1$, implying that PCL requires that a reduction in the bonus is compensated by a higher base salary. Thus, as long as PCL is satisfied, an altruistic principal always pays a higher base salary than an egoistic principal. The analysis thus leads to the following proposition:

Proposition 2.3 *In a separating equilibrium where the principal does not screen workers, an altruistic principal offers weaker financial incentives and pays a higher base salary than an egoistic principal.*

Given that the altruistic principal chooses the bonus that maximizes the expected surplus, how does he ensure that an egoistic principal is not willing to imitate him? This is the role of the base salary: the altruist increases the base salary up to the point that the selfish principal is not any longer willing to mimic him.

However, this does not imply that the altruistic principal pays a higher expected total compensation than his egoistic counterpart. The altruistic principal only pays a higher expected total compensation when workers provide more effort on average. To see this, it is convenient to use the fact that ICC1 is binding in equilibrium. Rewriting ICC1 gives:

$$b_h e_{\delta h} + s_h = b_l e_{\delta l} + s_l + e_{\delta h} - e_{\delta l}$$

Clearly, whether expected total compensation paid by the altruistic principal exceeds that of the selfish principal depends on the difference in average effort $e_{\delta h} - e_{\delta l}$. When the altruistic principal's equilibrium contract does not induce workers to provide more effort on average, credibly signaling altruism does not require paying a higher expected total compensation. As shown in the observable types case, conditionally altruistic workers do not always exert more effort when employed by an altruistic principal, because an altruistic principal sets weaker financial incentives, i.e. $b_h < b_l$. Therefore, we cannot be sure that a reciprocal worker exerts more effort when employed by an altruistic principal, whereas a selfish worker unambiguously provides less effort. Whether the altruistic principal pays a higher expected total compensation than an egoistic principal thus depends on the parameters.

Proposition 2.4 *The altruistic principal pays a higher expected total compensation than the egoist if and only if* $\gamma_h(\alpha_h) + r\theta\sigma^2 > 1$.

Proof. See appendix. ■

The intuition behind this condition is that the difference in average effort $e_{\delta h} - e_{\delta l}$ is only positive when the productivity of a conditionally altruistic worker sufficiently exceeds the productivity of an egoistic worker, either because a conditionally altruistic worker cares a lot about the principal's welfare or because it is costly to motivate the worker via financial incentives. So, combining propositions 3 and 4, paying a relatively low expected total compensation does not necessarily disprove a principal's altruism as long as it is accompanied by weak financial incentives.

The equilibrium is illustrated by figure 1. The figure shows ICC1 for $\gamma_h(\alpha_h) < \frac{1}{2}$ and the participation constraints PCL and PCH. Thus, ICC1 represents an isoprofit curve that for each bonus indicates the minimum base salary required to keep the selfish principal from imitating.[14] Similarly, PCL and PCH are indifference curves indicating the lowest base salary that is acceptable to a selfish and reciprocal worker, respectively. The arrows thus demonstrate the area of feasible contracts. The dotted line represents an indifference curve of a hypothetical 'average' worker. That is, it is a weighted average of both worker types' indifference curves representing the utility levels they obtain in the optimum. The optimum is where the indifference curve of the hypothetical 'average' worker is tangent to ICC1. The corresponding bonus (b_h) maximizes the expected surplus.

For the remainder of this chapter, it is important to understand the intuition behind the curves. The slope of ICC1 depends on $\gamma_h(\alpha_h)$:

Lemma 2.4 *Let s_h^{ICC1} denote the base salary that keeps the egoist from imitating, as defined by ICC1. s_h^{ICC1} is always decreasing in the bonus provided $(1-\delta)\gamma_h(\alpha_h) > \frac{1}{2}$, but has an inverted u-shape when $(1-\delta)\gamma_h(\alpha_h) < \frac{1}{2}$.*

Proof. See appendix. ■

The intuition is that ICC1 consists of two effects. On the one hand, worker's effort is increasing in the bonus, requiring an increase in the base salary to keep the selfish principal from mimicking. On the other hand, an increase in the bonus reduces the share of the marginal product that accrues to the principal, allowing for

[14] Specifically, ICC1 represents the isoprofit curve of the egoistic principal who imitates the altruist. This is identical to the material part of the altruistic principal's isopayoff curve.

a decrease in the base salary. Since highly reciprocal workers exert relatively high effort and are relatively insensitive to incentive pay, the latter effect dominates when $(1-\delta)\gamma_h(\alpha_h) > \frac{1}{2}$.

Both PCL and PCH slope downwards, as depicted in figure 1. The reason is that an increase in the expected bonus payment benefits the worker, implying that the base salary should decrease to keep expected utility constant. This effect is partially offset by an increase in the worker's exposure to risk.[15] The reciprocal worker's PC, PCH, is always below PCL, which follows from the fact that a reciprocal worker derives more utility from the same equilibrium contract than an egoist. Moreover, PCH has a steeper slope than PCL:

Lemma 2.5 *Let s_h^{PCL} and s_h^{PCH} denote the base salary such that PCL and PCH are exactly satisfied. s_h^{PCL} and s_h^{PCH} have the following properties:*

$$\frac{ds_h^{PCH}}{db_h} < \frac{ds_h^{PCL}}{db_h} < 0.$$

Proof. See appendix. ■

The reason is that a conditionally altruistic worker exerts more effort than an egoist for a given bonus, implying that an increase in the bonus leads to a larger increase in expected payments for a conditionally altruistic worker than for an egoistic worker. Therefore, a reciprocal worker permits a larger decrease in the base salary while keeping expected utility at the same level.

Two important observations need to be made. The first is that there is always a point on PCL that represents the contract offered by the selfish principal, namely (b_l, s_l). Since PCL is the selfish worker's indifference curve yielding his reservation utility $E(u_l) = \overline{u}$, the contract (b_l, s_l) is necessarily a point on PCL, as depicted in figures 1 and 2. The second observation is that when the altruistic principal offers b_h equal to b_l, ICC1 requires that $s_h > s_l$. The reason is that all contracts on ICC1 are assumed to succeed in signaling (and screening in the next section). Since for a given bonus expected effort is higher when the principal is believed to be an altruist, it must be that for $b_h = b_l$, $s_h > s_l$ to discourage the egoistic principal from imitating.

[15] This is by assumption, as I imposed that $r\sigma^2\theta \leq 1$ to ensure that indifference curves are downward sloping.

Finally, recall that I still have to show that the equilibrium contract (b_h, s_h) satisfies PCL. As argued above, the altruistic principal maximizes the utility of a hypothetical 'average worker', as workers are egoistic with probability δ. By lemma 5, the indifference curves of the 'average' worker are steeper than those of the egoist. Suppose that the altruistic principal chooses $b_h = b_l$ and $s_h > s_l$ to discourage the egoistic principal from imitating. This contract ensures that the egoistic worker's participation constraint is satisfied. Suppose that the altruistic principal chooses a contract (b_h, s_h), where $b_h < b_l$, on the 'average' worker's indifference curve that intersects this contract. As the average indifference curve is steeper than PCL, the selfish worker gains from a decrease in the bonus $(b_h < b_l)$ and a corresponding raise in the base salary that keeps the utility of the 'average' worker unchanged. Thus, the equilibrium contract ensures that the egoistic worker's participation constraint is satisfied.

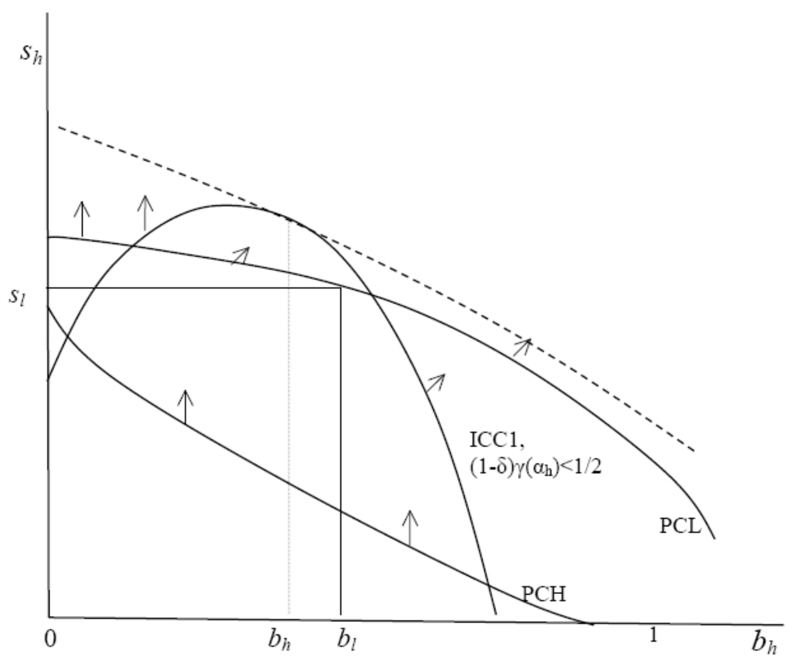

Figure 1: separating equilibrium without screening

Separating equilibrium: screening

This section studies a separating equilibrium where the altruistic principal writes a contract that signals his altruism and simultaneously screens worker types. First, consider the egoistic principal's contract choice. Because all workers behave completely egoistic when they believe that the principal is selfish (i.e. $\gamma_i(\alpha_l) = 0$), a selfish principal has no signaling or screening motive in equilibrium. Thus, as derived formally in lemma 1, the selfish principal offers a bonus $b_l = \frac{1}{1+r\sigma^2\theta}$ and the lowest base salary workers are willing to accept.

The altruistic principal may benefit from writing a contract that selects reciprocal workers only. As noted earlier, the altruistic principal's utility is increasing in the fraction of reciprocal workers, provided he convinces them that he is an altruist. Screening of worker types requires that the contract (b_h, s_h) simultaneously violates PCL and satisfies PCH. For ease of exposition, I refer to violating PCL as satisfying the screening constraint (SCC):

$$b_h e_{lh} + s_h - \frac{1}{2}\theta e_{lh}^2 - \frac{1}{2} r\sigma^2 b_h^2 \leq \overline{u}. \tag{SCC}$$

The proposed equilibrium contract should not only screen worker types, but also signal the principal's altruism. Thus, the contract should satisfy the two incentive compatibility constraints. Assuming that SCC and PCH are satisfied, the incentive compatibility constraints ICC1 and ICC2 can be written as:

$$(1 - b_l)e_{\delta l} - s_l \geq (1 - b_h)e_{hh} - s_h \tag{ICC1'}$$

$$(1 - b_l)e_{\delta l} - s_l + \alpha_h \overline{u} \leq (1 - b_h)e_{hh} - s_h + \alpha_h E(u_h) \tag{ICC2'}$$

Since ICC1' and ICC2' are nothing but special cases of ICC1 and ICC2, the same reasoning applies to show that ICC1' is always binding and hence ICC2' is slack. Because profits are increasing in the worker's altruism, a selfish principal is always willing to imitate an altruistic principal, unless the altruistic principal explicitly takes ICC1' into account, see lemma 3 for a formal proof. Moreover, we observed in the previous section that when the principal abstains from screening workers, PCL is always satisfied. Therefore, succesful screening requires that the principal takes SCC into account, in other words SCC is binding as well. This also proves that the reciprocal worker's PC is satisfied, because when SCC is binding, PCH is slack. Thus, the altruistic principal's contract offer can be characterized as follows:

Lemma 2.6 *In a separating and screening equilibrium, the altruistic principal offers a contract such that ICC1' and SCC hold with equality. Such a contract always exists for* $b_h \in (b_l, 1)$.

Proof. See appendix. ■

Figure 2 illustrates the altruistic principal's maximization problem. The similarity with figure 1 should be clear: it shows ICC1', SCC and PCH. Recall that the altruistic principal chooses the point on ICC1' that maximizes the reciprocal worker's expected utility. By shifting the conditionally altruistic worker's indifference curve (PCH) up along ICC1', it can easily be seen that, given the screening constraint, his expected utility is maximized at an intersection of ICC1' and SCC.

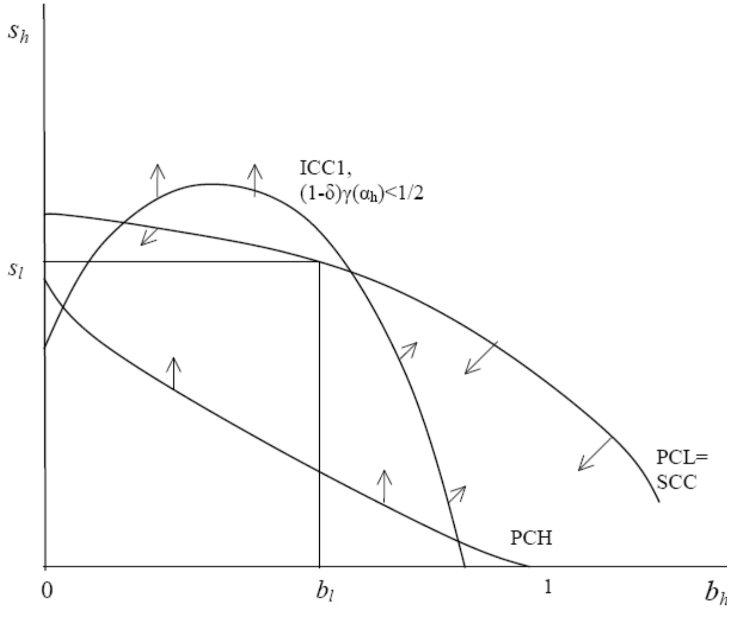

Figure 2: separating equilibrium where the altruistic principal screens.

Figure 2 provides an intuitive explanation of lemma 6. Recall that for $b_h = b_l$, ICC1' lies above SCC. When $b_h = 1$, ICC1' is always below SCC, and hence an intersection point on the interval $(b_l, 1)$ always exists. There is a clear intuition for this fact. When $b_h = 1$, the worker is the full residual claimant and SCC thus specifies the maximum amount he is willing to pay for the firm. This amount is equal to the expected revenues minus the costs of effort, risk-bearing and the outside option. Similarly, when $b_h = 1$, ICC1' specifies the maximum amount an altruistic principal can receive for the firm: this amount should not exceed the profits made by the selfish principal. The selfish principal's profits are given by the expected revenues minus the compensation for the worker's effort, risk and outside option. The amount the worker is willing to pay for the firm (SCC) is always smaller than the equilibrium profits made by the selfish principal (ICC1'), because the selfish principal sets the bonus at the surplus-maximizing level ($b_l < 1$). By contrast, when the altruistic principal makes the worker full residual claimant ($b_h = 1$), he exposes the worker to an inefficient amount of risk, which reduces his willingness to pay for ownership of the firm. When expressed as a base salary, SCC thus always exceeds ICC1', implying that an intersection point on the interval $(b_l, 1)$ always exists.

However, as illustrated by figure 2, when $\gamma_h(\alpha_h)$ is small, ICC1' has an inverted u-shape, implying that there are two intersection points.[16] In that case, the principal prefers the bonus at the intersection point on the interval $(b_l, 1)$. To see this, recall that the altruistic principal chooses the point on ICC1' that maximizes the reciprocal worker's expected utility. By shifting the conditionally altruistic worker's indifference curve (PCH) up along ICC1', it can easily be seen that, given the constraints, the principal prefers the intersection point that specifies the highest bonus. The reason is that an increase in the bonus is more beneficial for a reciprocal worker than an egoistic worker due to the former's higher effort.

Proposition 2.5 *In a separating and screening equilibrium, the altruistic principal offers stronger incentives and pays a lower base salary than the egoistic principal, i.e.* $b_h \in (b_l, 1)$ *and* $s_h < s_l$.

[16] A second intersection point may also exist when $\gamma_h(\alpha_h)$ is sufficiently high and $b_h > 1$. I assume that $b \leq 1$, but it can be shown that the intersection point on the interval $(b_l, 1)$ is strictly preferred.

Proof. See appendix. ■

In contrast to the previous section, the altruistic principal pays a lower base salary than the egoistic principal. As illustrated by figure 2, the downward sloping participation constraint of the egoistic worker implies that, as this constraint is binding, a higher bonus is reflected in a lower base salary. Nevertheless, the expected total compensation earned/paid by the altruistic worker/principal is always larger than that of the selfish type.[17] The reason is that because a reciprocal worker faces stronger incentives than a selfish worker, he unambiguously provides more effort ($e_{hh} > e_{\delta l}$), implying that an altruistic principal has to pay more than a selfish principal to discourage him from imitating. This result is summarized in the following proposition:

Proposition 2.6 *In a separating and screening equilibrium, an altruistic principal pays a higher expected total compensation than an egoistic principal.*

These results stand in remarkable contrast with the results in the previous section. The reason for these diverging findings is that screening can (most efficiently) be accomplished by offering stronger incentives than otherwise optimal. These excessively strong incentives reduce the total surplus and diminish the attractiveness of the contract for the selfish worker, which is inevitable in order to satisfy both the screening constraint and ICC1'.

Comparison of equilibria

One may wonder whether the altruistic principal prefers the separating and screening equilibrium above the standard separating equilibrium. This question is particularly relevant, because the altruistic principal can always deviate to the equilibrium that gives him the largest total expected payoff. As argued above, as long as ICC1 or ICC1' are satisfied, the intuitive criterion implies that deviation must come from

[17] As in the previous section, this can be seen from rewriting ICC1':

$$b_h e_{hh} + s_h = b_l e_{\delta l} + s_l + e_{hh} - e_{\delta l},$$

implying that $b_h e_{hh} + s_h > b_l e_{\delta l} + s_l$, since $e_{hh} > e_{\delta l}$.

an altruistic principal. Thus the equilibrium that yields the altruistic principal the highest payoff will be the unique equilibrium outcome.

Because the altruistic principal's profits are identical in the two situations (namely the same as the selfish principal's payoff $E(\pi_l^*)$), the equilibrium that yields the highest expected worker utility is preferred. Screening of worker types has the advantage that only conditionally altruistic workers are attracted, which has a positive effect on expected worker utility. First, because conditionally altruistic workers exert more effort, which is reflected in a higher salary. Second, because conditionally altruistic workers derive utility from the principal's welfare. However, screening is also costly: the bonus is distorted compared to the efficient bonus level. The higher $\gamma_h(\alpha_h)$, the larger the distortion. Thus, incurring the costs of screening is unattractive when $\gamma_h(\alpha_h)$ is relatively large and when the large majority of workers is conditionally altruistic:

Proposition 2.7 *A separating and screening equilibrium always exists for small values of $\gamma_h(\alpha_h)$. Screening is less likely, the higher $\gamma_h(\alpha_h)$ and the smaller δ.*

Proof. See appendix. ■

2.6 Concluding remarks

I have studied the relation between monetary gift-exchange and incentives by incorporating reciprocity in an otherwise standard principal-agent model. The specification of reciprocity is taken from Levine (1998), and allows to distinguish between authentic and strategic kindness. The key assumption is that some workers care more for the principal when they are convinced that the principal cares for them. The principal can be egoistic or altruistic. An altruistic principal can signal his type by offering a generous contract, consisting of a base salary and a piece rate. As is common in principal-agent models, the worker is risk-averse and not constrained by limited-liability. Inspired by the findings in several experiments, I have allowed for worker heterogeneity by assuming that not all workers are reciprocal. As a result, the principal may find it attractive to screen workers.

Assuming that types are private information, I have found that an altruistic principal who abstains from screening, signals his altruism by offering relatively weak incentives and a relatively high base salary. The piece rate and the base salary simultaneously convince the worker of the principal's care. The reason for offering weak incentives is that when workers are convinced of the principal's care, strong incentives add little to worker's productivity, while exposing workers to unnecessary risk. Offering strong incentives is therefore suboptimal for a principal who cares about the worker's well-being.

The second finding is that to induce reciprocity, an altruistic principal does not necessarily have to pay a higher expected total compensation than an egoistic principal. As part of the altruistic principal's gift is a reduction in incentives, the worker's expected effort may be relatively low, despite his altruistic feelings. Therefore, an egoistic principal has no reason to mimic the altruist, implying that a relatively low expected total compensation may suffice to distinguish both types. An altruistic principal only pays a higher expected total compensation when signaling altruism leads to significant productivity gains.

Finally, as some workers do not reciprocate the principal's altruism, the principal may find it optimal to write a contract that simultaneously signals his altruism and screens reciprocal worker types. I have shown that such a contract is characterised by excessively strong incentives and a relatively high expected total compensation. Incentives are a suitable instrument for screening workers, because conditionally altruistic workers put in more effort than egoistic workers and hence gain more from output-contingent pay. Thus, strong incentives are offered to attract the worker who needs them the least.

Of course, there are some limitations to the analysis. A first limitation is that I only looked at monetary rewards, while employers typically have other instruments to stimulate or control workers, such as work rules, work organisation, minimum effort requirements, task assignment or giving personal attention to workers. All of these may be helpful to signal the employer's benevolent intentions. Therefore, it may well be that employers that establish good relationships with their workers via these non-monetary means can afford paying lower wages, as in Dur (2009) and Dur et al. (2010). This would be well in line with the results of the observable types case.

Relatedly, a second limitation is that in large organizations wages are not always determined by the relevant managers, and the interpretation as a gift may therefore be problematic. It would be interesting to see how wage-setting institutions and organizational structure impact on the prospects for gift-exchange.

A third limitation is that the model I presented is a partial equilibrium model: there is no competition between employers and the outside option of both worker types is exogenously given. It is not ex-ante clear to what extent the results carry over to a general equilibrium setting. Thus, there is ample room for further research, both theoretically and empirically.

2.A Appendix

Proof of Proposition 2

The proof builds on the analysis of the separating equilibria. Screening requires that the egoistic worker's participation constraint (PCL) is violated, while at the same time the conditionally altruistic worker's participation constraint is satisfied. All contracts that satisfy these two conditions are possible equilibrium candidates, depending on the out-of-equilibrium beliefs. I show that none of these contracts, denoted (b_p, s_p), is an equilibrium contract, as there is always a profitable deviation.

First, we can rule out all pooling contracts that yield the egoistic principal lower expected profits than he earns in a separating equilibrium, $E(\pi_l^*)$. The egoistic principal can always profitably deviate by offering the separating equilibrium contract (b_l, s_l), which ensures him at least $E(\pi_l^*)$, irrespective of the worker's out-of-equilibrium beliefs.

To rule out all pooling contracts (b_p, s_p) that yield higher profits than $E(\pi_l^*)$, we proceed along the lines suggested in the main text. We look for an alternative contract offer (b', s') that can be profitable for an altruistic principal, but not for an egoist. The reason is that for such a deviation, the Intuitive Criterion uniquely defines the out-of-equilibrium beliefs, namely that $pr(\alpha_j = \alpha_h) = 1$. More formally, the contract (b', s') breaks the equilibrium if the following conditions are satisfied:

1. The expected monetary payoff with the most optimistic beliefs possible is equal to the monetary payoff in equilibrium.
2. A conditionally altruistic worker obtains higher expected utility than when the equilibrium contract (b_p, s_p) is offered.
3. The screening contraint SCC is satisfied, i.e. the egoistic worker's participation constraint PCL is violated.

That such a contract always exists can be shown as follows. Let ICC1" represent the same isoprofit curve as ICC1', but for a higher level of profits, namely $E(\pi_p) > E(\pi_l^*)$. Mathematically:

$$(1 - b')e_{hh} - s' = E(\pi_p), \tag{ICC1"}$$

Thus, ICC1" represents all contracts (b', s') that yield the same profit level $E(\pi_p)$

as the equilibrium contract (b_p, s_p), assuming the worker is conditionally altruistic and that he believes that he is employed by an altruistic principal. As higher profit levels are represented by lower indifference curves, all pooling contracts that yield higher monetary profits than (b_l, s_l) have a corresponding isoprofit curve ICC1" below ICC1'.

Suppose the altruistic principal increases the base salary up to the point that monetary profits equal equilibrium profits, i.e. $b' = b_p$, $s' > s_p$ such that ICC1" holds. There are two cases. First, the screening constraint SCC is satisfied. Clearly, such a contract meets all of the conditions for profitable deviation. Second, the screening constraint SCC is violated. In that case the isoprofit curve ICC1" always contains another contract that meets the conditions outlined above, characterised by $b' \in (b_p, 1)$. As ICC1" is nothing but ICC1' for a higher profit level, it follows from the properties of ICC1' that for $b' = b_p$, $s' > s_p$, while for $b' = 1$ PCL is violated (or SCC satisfied). Thus, in this case there always exists a contract $b' \in (b_p, 1)$ and corresponding base salary s' such that PCL and ICC1" are exactly satisfied. This contract therefore also yields a conditionally altruistic worker higher utility: increasing the base salary until PCL is satisfied, while keeping b constant, gives a conditionally altruistic worker a higher utility than in equilibrium. As the participation constraint of a conditionally altruistic worker is steeper than PCL, offering $b' \in (b_p, 1)$ further increases his utility.

Proof of proposition 4

As ICC1 is binding, the altruistic principal pays a higher expected total compensation if and only if $e_{\delta h} > e_{\delta l}$. Inserting expressions for effort, b_l and b_h into this condition, we obtain after considerable rewriting:

$$\frac{(1-\delta)\, r\sigma^2 \gamma_h(\alpha_h)}{(1+\theta r\sigma^2)} \frac{\gamma_h(\alpha_h) + r\theta\sigma^2 - 1}{\delta + (1-\delta)\left[1-\gamma_h(\alpha_h)\right]^2 + \theta r\sigma^2} > 0,$$

implying that the altruist pays more if and only if

$$\gamma_h(\alpha_h) + r\theta\sigma^2 > 1.$$

Proof of lemma 4

It follows from straightforward rewriting of ICC1 that s_h^{ICC1} is defined as

$$s_h^{ICC1} = (1-b_h)e_{\delta h} - (1-b_l)e_{\delta l} - s_l,$$

or rewritten, using equations (2.5) and (2.6):

$$s_h^{ICC1} = (1-b_h)\left(\frac{b_h}{\theta} + (1-\delta)\,\frac{\gamma_h(\alpha_h)(1-b_h)}{\theta}\right) - \frac{1}{2\theta}\frac{1}{(1+\theta r\sigma^2)} + \overline{u}.$$

Inspection of the derivative with respect to b_h proves that s_h^{ICC1} initially increases in the bonus provided $(1-\delta)\,\gamma_h(\alpha_h) < \frac{1}{2}$, but always decreases in the bonus when $(1-\delta)\,\gamma_h(\alpha_h) > \frac{1}{2}$:

$$\frac{ds_h^{ICC1}}{db_h} = \frac{-2\left[b_h + (1-\delta)\,\gamma_h(\alpha_h)(1-b_h)\right] + 1}{\theta}.$$

Proof of lemma 5

Assuming that PCL or, equivalently, SCC holds with equality, PCL can be rewritten to:

$$s_h^{PCL} = \overline{u} - \frac{b_h^2}{2\theta}\left(1-\theta r\sigma^2\right).$$

Clearly, s_h^{PCL} decreases in the bonus by assumption, as we imposed that $\theta r\sigma^2 < 1$ to prevent un upward slope. For ease of comparison, I provide the derivative with respect to b_h:

$$\frac{ds_h^{PCL}}{db_h} = -\frac{b_h}{\theta}\left(1-\theta r\sigma^2\right).$$

Similarly, s_h^{PCH} follows from rewriting PCH:

$$s_h^{PCH} = \overline{u} - \frac{b_h^2}{2\theta}\left(1-\theta r\sigma^2\right) + \frac{\gamma_h(\alpha_h)^2}{2\theta}\left(1-b_h\right)^2 - \gamma_h(\alpha_h)E(\pi_h).$$

Keeping $E(\pi_h)$ constant at its equilibrium level $E\left(\pi_l^*\right)$ and taking the derivative with respect to b_h, it can easily be seen that compared to PCL, PCH is steeper as long as $b_h \leq 1$:

$$\frac{ds_h^{PCH}}{db_h} = -\frac{b_h}{\theta}\left[1-\theta r\sigma^2\right] - \frac{\gamma_h(\alpha_h)^2}{\theta}\left[1-b_h\right].$$

Proof of lemma 6

Because both ICC1' and SCC are binding, the altruistic principal maximizes his payoff by offering a contract such that both constraints hold with equality. To show that such a contract always exists, first note that for $b_h = b_l$, ICC1' requires a higher base salary s_h than SCC. Since for a given bonus, a conditionally altruistic worker exerts more effort when he believes that the principal is an altruist, signaling is only credible if $s_h > s_l$. It can be shown that when $b_h = 1$, ICC1' always allows for a lower base salary than SCC, and hence an intersection point on the interval $(b_l, 1)$ always exists.

Recall that s_h^{PCL} denotes the maximum salary SCC allows for, see lemma 5. Inserting $b_h = 1$ into SCC, $-s_h^{PCL}$ describes a selfish worker's willingness to pay for the firm:

$$-s_h^{PCL} = e_{lh} - \frac{1}{2}\theta e_{lh}^2 - \frac{1}{2}r\sigma^2 - \overline{u}.$$

Recall that s_h^{ICC1} denotes the base salary that keeps the egoistic principal from imitating, see lemma 4. As ICC1' is nothing but a special case of ICC1, $s_h^{ICC1'}$ can be obtained by inserting $\delta = 0$ into s_h^{ICC1}. Inserting $b_h = 1$ into ICC1', we obtain that $-s_h^{ICC1'} = E(\pi_l^*)$. Thus, the maximum amount the altruistic principal can receive for the firm $(-s_h^{ICC1'})$ is equal to the egoistic principal's expected equilibrium profits. Using equations (2.8) and (2.10), we can write $E(\pi_l^*)$ as:

$$-s_h^{ICC1'} = E(\pi_l^*) = e_{ll} - \frac{1}{2}\theta e_{ll}^2 - \frac{1}{2}r\sigma^2 b_l^2 - \overline{u}.$$

Since b_l is the surplus maximizing bonus when workers are egoistic, $E(\pi_l^*) > -s_h^{PCL}$, implying that $s_h^{ICC1'} < s_h^{PCL}$ when $b_h = 1$.

Proof of proposition 5

By lemma 6, the equilibrium contract satisfies ICC1' and SCC with equality, and such a contract always exists for $b_h \in (b_l, 1)$. The altruistic principal prefers this contract over possible contracts that also satisfy ICC1' and SCC with equality for $b_h < b_l$. As the principal's monetary payoff is constrained by ICC1', he chooses the contract that maximizes a conditionally altruistic worker's utility. By lemma 5, a conditionally altruistic worker's indifference curve has a steeper slope than an egoistic worker's indifference curve. Hence, utility of a conditionally altruistic worker is maximized when the principal chooses $b_h \in (b_l, 1)$ rather than $b_h < b_l$. The downward sloping indifference curves imply that $s_h < s_l$.

Proof of proposition 7

To show under what conditions the altruistic principal prefers to screen, first note that when $\gamma_h(\alpha_h) = 0$, the payoff is equal in both equilibria. I now analyze how the principal's expected payoff changes when $\gamma_h(\alpha_h)$ increases in both equilibria, where the change in the principal's expected payoff is equal to the change in the worker's expected utility.

In a separating equilibrium where the principal does not screen, the effect of a change in $\gamma_h(\alpha_h)$ on total utility is given by $\delta \frac{dE(u_l)}{d\gamma_h(\alpha_h)} + (1-\delta)\frac{dE(u_h)}{d\gamma_h(\alpha_h)}$, or after some rewriting:

$$\delta \frac{dE(u_l)}{d\gamma_h(\alpha_h)} + (1-\delta)\frac{dE(u_h)}{d\gamma_h(\alpha_h)} = (1-\delta)\left[(1-b_h)\frac{de_{hh}}{d\gamma_h(\alpha_h)} + E(\pi_h)\right].$$

This has a simple interpretation; the gain in total utility when $\gamma_h(\alpha_h)$ increases is equal to the additional productivity of the reciprocal type (reflected in the base salary), plus his increased utility from the immaterial aspect of the job. Worker's utility is convex in $\gamma_h(\alpha_h)$, because b_h is decreasing and hence $\frac{de_{hh}}{d\gamma_h(\alpha_h)}$ is increasing in $\gamma_h(\alpha_h)$.

Similarly, in a separating and screening equilibrium the effect of a change in $\gamma_h(\alpha_h)$ on total expected utility is given by $\frac{dE(u_h)}{d\gamma_h(\alpha_h)}$, or:

$$\frac{dE(u_h)}{d\gamma_h(\alpha_h)} = \frac{db_h}{d\gamma_h(\alpha_h)}\left[\frac{de_{hh}}{db_h}(1-\theta e_{hh}) - r\sigma^2 b_h\right] + (1-b_h)\frac{de_{hh}}{d\gamma_h(\alpha_h)} + E(\pi_h).$$

The first part is negative and represents the loss in worker utility because b_h is suboptimally high. The second part has a similar interpretation as in a signaling equilibrium, but keeping the bonus constant, it is larger because the worker is reciprocal with probability 1 instead of $(1-\delta)$. The first part is zero if $\gamma_h(\alpha_h) = 0$, but becomes smaller (increases in absolute value) when $\gamma_h(\alpha_h)$ becomes larger. The second part is concave in $\gamma_h(\alpha_h)$, because b_h is increasing and hence $\frac{de_{hh}}{d\gamma_h(\alpha_h)}$ is decreasing in $\gamma_h(\alpha_h)$.

Comparing $\delta \frac{dE(u_l)}{d\gamma_h(\alpha_h)} + (1-\delta) \frac{dE(u_h)}{d\gamma_h(\alpha_h)}$ and $\frac{dE(u_h)}{d\gamma_h(\alpha_h)}$, the separating and screening equilibrium is always preferred for $\gamma_h(\alpha_h)$ sufficiently close to zero. When $\gamma_h(\alpha_h)$ increases, the cost of distorting the bonus becomes more severe, and at some point the standard separating equilibrium will be preferred. The larger the fraction of selfish workers δ, the higher the values of $\gamma_h(\alpha_h)$ that sustain a screening equilibrium.

Chapter 3

Reciprocity and Incentive Pay in the Workplace

Joint with Robert Dur and Hein Roelfsema[1]

3.1 Introduction

The birth of modern management theory is often related to the Hawthorne studies (Roethlisberger and Dickson 1939). A well-known case among social scientists, in the Hawthorne plant researchers experimented with light intensity and examined its effect on worker's productivity. To their surprise, workers reacted favorably to both increases and decreases of light intensity. The conclusion of the researchers was that workers simply liked the attention of management and responded by increasing effort. This may come as no surprise to most organizational scholars today, but in the times of Taylorian scientific management with its job specialization and monetary incentives, such findings caused heated debate.

To a large extent, the debate on the importance of monetary incentives and other, non-monetary tools of management in motivating workers continues today. In economics, agency theory is characterized by a strong emphasis on monetary incentives. This is not without reason: the importance of monetary incentives for

[1] A slightly adapted version of this chapter is published in Journal of Economic Psychology, August 2010, vol. 31(4), pp. 676-686.

workers' motivation is confirmed in many empirical studies, see Prendergast (1999) for an overview. Strong motivational effects of monetary incentives are also found in recent field experiments (see, among others, Lazear 2000 and Shearer 2004).

However, monetary incentives are often not considered as the most important motivator by workers and managers. Many workers consider task enjoyment and moral concerns as stronger motivators than monetary incentives (Minkler 2004). In a study on managers' use of motivational tools, Agell (2004) reports that more than 60% of managers in Sweden use 'good management-worker relations' to a great or fairly great extent. In contrast, even in the sector where performance-related pay is most common (skilled services), only 17% of managers use performance pay as a motivational tool to a great or fairly great extent. Likewise, Campbell and Kamlani (1997) find that compensation executives in the US rank good management-worker relationships much higher than wages, working conditions, and supervision as determinants of workers' effort (see also Bewley 1999).

In line with these surveys, studies in organizational behavior (OB) stress the importance of so-called leader-member exchange relationships (LMX) and perceived supervisor support (PSS) for workers' motivation. A large number of empirical studies find a positive effect of the perceived quality of the management-employee relation on employee's performance, e.g. Nagin et al. (2002), Graafland and Rutten (2004), Shanock and Eisenberger (2006), Kamdar and Van Dyne (2007), Freeman et al. (2008), and Pazy and Ganzach (2009). Moreover, several studies find that better relations between managers and employees lead to higher job satisfaction and/or reduced turnover intentions, see e.g. Babin and Boles (1996) and Wayne et al. (1997). In line with these findings, empirical work also suggests that firms with bad management-worker relationships are penalized in that they need to pay higher wages to attract and retain workers (Pfeffer 1998, Gittell 2003, and Borzaga and Depedri 2005).

This chapter reconciles these two views on workers' motivation by developing a formal agency model that incorporates the OB tools of management and monetary incentives. We picture a firm where workers exert effort and managers, in addition to incentive pay, use non-monetary tools of management (attention, praise, recognition, et cetera). We make two important assumptions about these non-monetary tools, which are inspired by the empirical studies mentioned above. First, applying these tools of management – which we call management attention – raises worker's well-being but comes at a cost for the manager. Second, we assume that management attention reduces worker's marginal cost of effort, implying that effort increases with attention. In our model, the reason is worker's reciprocity: workers reciprocate management attention by providing effort. In this setting, we study the optimal provision of incentive pay for workers, the manager's incentive to apply non-monetary tools of management, and the resulting worker's behavior and productivity.

One of our key objects of study is the issue of 'congruence', important in management science, but not often studied in organizational economics. The idea is that the set up of one element of organization affects the working of other parts (see, e.g., Nadler and Tushman 1997). We argue that the strength of monetary incentives given to workers affects the extent to which managers use their other motivational tools. In particular, we will show that, when neither worker's effort nor management activities are contractible, incentive pay for workers weakens the incentive for managers to motivate workers through attention. The reason is that, by leaving a larger share of output to the worker, there is less to gain from increases in output for the manager, which dilutes his incentives to support the worker. Optimal performance pay for the worker therefore strikes a balance between motivating the worker to exert effort and preserving incentives for the manager to apply his non-monetary management tools. Our analysis thus predicts that managers will be careful with introducing or raising incentive pay for workers, and particularly so for workers who are most responsive to management attention, that is, workers who are highly reciprocal. In equilibrium, worker's effort and manager's attention are both suboptimally low compared to the first-best.

These results change when the manager employs multiple workers doing comparable tasks. Following Carmichael (1983)'s analysis of the 'agent-agents problem,' we show that first-best profits can then be achieved through promotion incentives for workers. The reason is that, in contrast to individual performance pay, promotion incentives do not interfere with the manager's incentive to give attention, because the total amount of wage compensation to the workers is fixed in advance. This benefit of promotion incentives is particularly large when workers are highly responsive to manager's attention.

The main predictions of our theoretical analysis are thus twofold: workers who are more responsive to manager's attention are less likely to receive individual performance pay and more likely to receive promotion incentives. We empirically examine these predictions using data from the German Socio-Economic Panel (GSOEP), which contains data on compensation schemes and reciprocity for more than 2700 German workers. While we find no support for the former hypothesis, there is strong support for the latter: Worker's reciprocity significantly increases the likelihood of receiving promotion incentives.

We proceed as follows. The following section gives a brief overview of related literature. Next we introduce in section 3.3 our basic model. Section 3.4 and 3.5 analyze optimal contracts and the resulting manager's and worker's behavior in the first-best, the second-best with full rent extraction and with limited-liability protection, respectively. Section 3.6 extends the analysis to allow for multiple workers and promotion incentives. Section 3.7 describes the results of our empirical analysis. Section 3.8 concludes.

3.2 Related literature

The economic literature on manager-subordinate reciprocity has so far been confined to monetary gift-exchange. Starting with Akerlof (1982), economists have argued that paying generous wages may trigger effort and loyalty as workers feel a need to reciprocate the employer's gift. Numerous laboratory experiments have provided support for this monetary gift-exchange relation (an early study is Fehr, Kirchsteiger, and Riedl 1993; Fehr and Gächter 2000 provide an overview of the voluminous literature). Recent field studies, however, are less supportive. In various natural workplace settings, Gneezy and List (2006), Kube et al. (2008), and Hennig-Schmidt et al. (2010) find only limited support for monetary gift-exchange.

As discussed in the Introduction, studies in management and organizational psychology have emphasized other managerial tools facilitating relationship-building between managers and employees, namely offering socioemotional resources that address "social and esteem needs (and are often symbolic and particularistic)." (Cropanzano and Mitchell 2005, p. 881). The main contribution of this chapter is to incorporate such social exchange as a management tool into an otherwise standard agency model, which allows us to study social exchange and several forms of incentive pay in one unifying framework.[2]

Closest to our study is Englmaier and Leider (2008)'s recent paper on the implications of reciprocity for the employment relation. Their key result is that reciprocal motivations and performance-based pay are substitutes, as in section 3.5 of this chapter. Their analysis strongly differs from ours, however. The crucial difference is that they confine their analysis to monetary gift-exchange, whereas we focus on social gift-exchange. Specifically, in their model, the principal is a passive contract-writer, inducing reciprocity by leaving a rent for the agent. In our model, feelings of reciprocity are engendered by the principal's attention. Further, we have a different approach in the empirical verification of our results. They provide empirical support by comparing the organizational form and pay structure between firms who supposedly select workers on reciprocity and those that do not. By contrast, we use the individual worker as the unit of analysis, using a direct measure of an individual's reciprocity.

This chapter builds on a rich body of literature that studies optimal contracts in the presence of double moral-hazard (e.g. Carmichael 1983, Demski and Sappington 1991, and Gupta and Romano 1998). Applications in the field of franchising are especially interesting, because some serious efforts have been made to empirically verify the theoretical predictions. For instance, Lafontaine (1992) finds that franchising contracts "are most consistent with a model based on two sided-moral hazard" (p. 263). Agrawal and Lal (1995) find "support for the hypothesis that the royalty rate balances the incentives to the franchisor to invest in brand name with those to the franchisees to invest in retail service" (p. 213). We differ from this literature in our focus on social exchange and workers' motivation. Further, we are

[2] In a related paper, we study social exchange and incentive provision in a common agency context, see Dur and Roelfsema (2010).

the first to provide some (indirect) evidence for the relevance of double moral-hazard in the workplace.

3.3 The model

We consider a risk-neutral principal employing a risk-neutral agent. The principal's expected payoff $E(\pi)$ is described by:

$$E(\pi) = eH + (1 - e)L - (w + eb) - \frac{1}{2}\rho a^2, \tag{3.1}$$

where e is effort exerted by the agent, H and L are the two possible output values (high and low; $H > L$), w is the agent's base salary, b is a bonus paid to the agent in case output is high, $\rho > 0$ is a cost parameter, and a denotes the principal's attention given to the agent. The probability that output is high is increasing in the agent's effort and, for simplicity, given by e where $e \in [0, 1]$. Throughout, we assume that the parameters are such that we can rule out solutions where $e < 0$ or $e > 1$ (see the working paper version of this chapter, Dur et al. (2008), for the exact conditions). Besides offering a contract describing the agent's base salary and bonus pay, the principal engages in giving attention to the agent. We assume that giving attention is costly for the principal. Allowing for some intrinsic benefits from giving attention would not change the results qualitatively.

The agent's expected utility $E(U)$ is:

$$E(U) = w + eb + \gamma ea - \frac{1}{2}\theta e^2, \tag{3.2}$$

where the first two terms are the agent's expected wage income and the last term represents the agent's cost of effort, $\theta > 0$. A distinguishing feature of our model is the interaction term γea, where $\gamma \geq 0$. This term captures the observations discussed in the Introduction that the agent's marginal costs of effort are decreasing and the worker's well-being is increasing in the attention given by the principal. Attention can be interpreted as kindness that evokes feelings of reciprocity, but it can also be interpreted as support by the principal which helps the agent to perform his tasks. We shall call γ the agent's reciprocity parameter.[3]

The timing of the game is as follows. First, the principal writes a contract which the agent accepts if his expected utility is equal to or exceeds his reservation utility $\overline{U}$. In the second stage the principal decides on his attention level. Finally, the agent chooses his effort level, after observing the attention provided by the principal.

3.4 First-best contract

Let us first consider the benchmark case where both attention and effort are contractible. Full contractibility implies that there is no reason to pay a bonus conditional on output on top of the base salary, and so we impose $b = 0$ in this section. Thus, the principal designs a contract consisting of attention a, effort e, and base salary w, that maximizes his expected payoff given by (3.1), taking into account the agent's participation constraint $E(U) \geq \overline{U}$, where $E(U)$ is given by (3.2). After rewriting the first-order conditions for attention and effort, we obtain:

$$a = \frac{\gamma e}{\rho}, \tag{3.3}$$

$$e = \frac{H - L + \gamma a}{\theta}. \tag{3.4}$$

Note that attention is increasing in effort and vice versa. In other words, attention and effort are complements. Thus, a change in ρ, θ, $(H - L)$, or γ affects the optimal values of both attention and effort. For example, when the principal's cost of giving attention ρ increases, it is optimal to pay less attention to the agent and, as a consequence of the higher marginal cost of effort, agent's optimal effort decreases.

3.5 Incomplete contract with full rent extraction

Next consider the situation where neither attention nor effort are contractible. Hence, in order to induce the agent to exert effort, the principal may find it op-

[3] Alternatively, we could assume that the principal's attention affects an agent's degree of altruism towards the principal, thus inducing the agent to exert more effort in response to attention as long as the agent is not the full residual claimant. See Dur (2009) for a model along these lines.

timal to offer bonus pay conditional on high output in addition to the base salary. We assume throughout that the agent's participation constraint, $E(U) \geq \overline{U}$, is binding. However, we will briefly discuss what happens when the agent is protected by limited liability at the end of this section. We solve for a subgame perfect equilibrium by backward induction. So, we start with the agent's effort choice, next we study the principal's choice of attention, and finally we consider optimal contract design.

The agent maximizes his expected utility (3.2) by choosing an optimal effort level, given the level of attention and the wage contract. The first-order condition for optimal effort implies:

$$e = \frac{b + \gamma a}{\theta}. \tag{3.5}$$

Obviously, the higher the bonus and the lower the costs of exerting effort, the higher the agent's effort. Further, effort increases with attention. Comparing (3.5) with (3.4) gives the usual result that, for a given a, the agent chooses first-best effort when the bonus equals the full marginal product ($b = H - L$).

The principal chooses attention so as to maximize his expected profits (3.1) taking the agent's reaction function (3.5) into account. The first-order condition is:

$$\frac{de}{da}(H - L - b) - \rho a = 0. \tag{3.6}$$

Rewriting using (3.5) gives optimal attention:

$$a = \frac{\gamma}{\theta\rho}(H - L - b), \tag{3.7}$$

which increases with the marginal product and with the agent's reciprocity, and decreases with the cost parameters. Moreover, equation (3.7) shows a clear negative relation between attention and the agent's bonus pay. The intuition for this result can be seen from the first-order condition (3.6). In the second stage of the game, the principal's only reason for giving attention is to stimulate effort. An increase in the bonus b reduces the principal's marginal payoff from the agent's effort, and hence reduces his optimal attention. In the extreme case that $b = H - L$, the full marginal product from effort accrues to the agent. In this case, no attention will be given, because there is nothing at stake for the principal. Another extreme case is $b = 0$, in which case all gains from extra effort accrue to the principal, and so giving attention is very attractive. Still, however, attention is below the first-best level.[4] The reason is that, after the contract has been signed, the principal only takes his own welfare into account and does not care about the positive effect his attention has on the agent's utility. To reach first-best attention, the bonus should actually be negative.

Anticipating the behavior of the principal and the agent in stage two and three of the game as described by equations (3.5) and (3.7), the principal writes a profit-maximizing contract in the first stage by choosing w and b, taking into account the agent's participation constraint $E(U) \geq \overline{U}$. After some rewriting we obtain the following first-order condition describing optimal bonus pay:

$$\frac{de}{db}(H - L - b) + \frac{da}{db}\gamma e = 0, \tag{3.8}$$

where, using (3.5) and (3.7), $de/db = 1/\theta$ and $da/db = -\gamma/\theta\rho$. The first-order condition clearly reveals the trade-off the principal faces when writing the contract. The bonus has a positive effect on effort, which increases profits as long as $b < H - L$.[5] On the other hand, the bonus reduces the amount of attention, which reduces the agent's utility by γe, and so increases total wage compensation. As a result, the optimal bonus is smaller than the marginal product if $\gamma > 0$. Using (3.5) and (3.7), it follows from (3.8) that the optimal bonus is given by:

$$b^* = \frac{\left(\theta^2\rho^2 - \gamma^4\right)(H - L)}{\theta^2\rho^2 - \gamma^4 + \gamma^2\theta\rho}, \tag{3.9}$$

which decreases in the agent's reciprocity γ for two reasons. First, when γ is higher, attention is more valuable to the agent. Second, when γ is higher, the principal's

[4] This is easily verified by solving (3.3) and (3.4) for a, and comparing with (3.7), substituting $b = 0$.

[5] As we have seen, a higher bonus implies higher effort for a given level of attention (see equation (3.5)), but leads to a reduction in attention (see equation (3.7)), which in turn reduces effort. Overall, effort increases in the bonus. This can be easily verified by substituting (3.7) into (3.5) and noting that $\gamma < \sqrt{\theta\rho}$, which ensures finite attention and effort in the first-best.

attention is more responsive to changes in the bonus. It is easily verified that the resulting levels of attention and effort are below their first-best levels. Clearly, the problem of having only one instrument (the bonus) for two conflicting goals (incentivizing the agent and the principal) implies that both attention and effort are suboptimally low.

Finally, consider the case where the agent is protected by limited liability: the agent's compensation must at least be equal to $\overline{w}$.[6] The results for the third and second stage of the game are identical, but the outcome of the contracting stage differs. In the full-rent-extraction case, the binding participation constraint makes it optimal for the principal to take the agent's welfare into account in the contracting stage, because the agent's welfare is reflected in the base salary. By contrast, when the limited-liability constraint is binding, bonus pay and attention do not reduce the base salary (which is fixed at $\overline{w}$), but increase the agent's rent. Therefore, the principal faces the standard trade-off between stimulating effort and leaving a rent to the agent.[7] Again, the bonus is decreasing in the agent's reciprocity γ, but for a different reason than in the full-rent-extraction case. When the agent is more reciprocally inclined, effort is higher for given values of attention and bonus pay. Therefore, the probability that the bonus actually has to be paid is higher, implying that providing a bonus is a more costly instrument when workers are more reciprocal (see Besley and Ghatak 2005 for a similar argument in the context of motivated workers). Hence, the optimal bonus decreases in agent's reciprocity.

3.6 Promotion incentives

This section examines the possibility of overcoming the double moral-hazard problem identified in the previous section through provision of promotion incentives (or relative performance pay).[8] Clearly, for this to be feasible, the principal must em-

[6] For brevity, we abstract from the case where both the participation constraint and the limited-liability constraint are binding. Our main results carry over to this case, however.

[7] The first-order condition describing optimal bonus pay is: $\frac{de}{db}(H - L - b) - e = 0$.

[8] An alternative way to overcome the double moral-hazard problem is to hire a middle-manager and to provide incentives so that he gives optimal attention to the worker. This, however, is more costly than promotion incentives, since the middle-manager needs to be compensated for his outside option. Option contracts a la Nöldeke and Schmidt (1995) do not solve the double moral-hazard

ploy at least two agents. For convenience, suppose the principal hires two identical agents, denoted by index numbers 1 and 2, who perform identical tasks.[9] The agents compete for a single promotion prize, which is denoted by P. We assume that the agent who achieves highest output wins the prize P. In case of equal outputs, a random draw determines the winner. Thus, the probability of winning the prize for agent 1 is given by $\frac{1}{2}(1 + e_1 - e_2)$, where the subscripts denote effort provided by the indicated agent. Expected utility for agent 1 is described by:

$$E(U_1) = w + e_1 b + \frac{1}{2}(1 + e_1 - e_2) P + \gamma e_1 a_1 - \frac{1}{2}\theta e_1^2. \tag{3.10}$$

The principal's payoff is described by

$$E(\pi) = (H - L - b)(e_1 + e_2) + 2L - (2w + P) - \frac{1}{2}\rho(a_1 + a_2)^2. \tag{3.11}$$

Note that the cost of the promotion prize does not depend on effort, because the principal awards the prize to one of the agents independent of the levels of output.

The analysis proceeds in the same way as in the previous section. Optimal third-stage behavior follows from the maximization of the agent's utility function (3.10) which results in:

$$e_1 = \frac{b + \frac{1}{2}P + \gamma a_1}{\theta}; \quad e_2 = \frac{b + \frac{1}{2}P + \gamma a_2}{\theta}. \tag{3.12}$$

These expressions are similar to our earlier findings (see equation (3.5)); the difference is that the agent is now also motivated by the possibility of winning the promotion prize P.

Optimal behavior in the second stage follows from maximization of the principal's payoff function (3.11) with respect to a_1 and a_2. Assuming that the principal gives the same level of attention to each of the two agents,[10] optimal attention provision

problem, because we deal with a case of two-sided direct externalities.

[9]The assumption of homogeneous agents is critical for the efficiency of promotion incentives. Kosfeld and von Siemens (2009) provide an argument for why workers with similar social preferences may sort into the same firms.

[10]It is easy to verify that in our set-up the distribution of a given level of attention over the agents does not influence agents' total effort and, hence, the principal's profits in the second stage of the game. Obviously, if agent's responsiveness to attention would decrease in the level of attention, it would be profit-maximizing to distribute attention evenly, as we assume.

is given by:

$$a_1 = a_2 = \frac{\gamma}{2\theta\rho}(H - L - b). \tag{3.13}$$

Equations (3.12) and (3.13) already make clear that the double moral-hazard problem can be solved by using promotion incentives. The promotion prize P incentivizes the agents, but does not impair the principal's incentives to give attention. Therefore, it is possible to set the bonus b such that the principal's incentives to provide attention are optimal, whereas the promotion prize P provides the agents with incentives to put in effort. Solving for the optimal contract, we obtain:

$$b^* = \left(1 - \frac{2\theta\rho}{2\theta\rho - \gamma^2}\right)(H - L),$$

$$P^* = \frac{4\theta\rho}{2\theta\rho - \gamma^2}(H - L).$$

Clearly, when $\gamma > 0$, it is optimal for the principal to offer promotion incentives to the agents along with negative individual bonus pay. Given our previous discussion in section 3.5, this result is not surprising. In the second stage, the principal does not take the agent's welfare into account. Therefore, in order to internalize this externality, the principal sets a negative bonus. This obviously gives perverse incentives to the agents, but this is corrected for by offering the promotion prize. By substituting the expressions for the optimal bonus and promotion prize into the equations for effort and attention (equations (3.12) and (3.13) respectively), it is easily verified that the first-best is reached.[11] As before, the bonus is decreasing in γ. The reason is that the size of the externality increases in γ, necessitating a lower bonus to reach first-best attention. Consequently, promotion incentives are also increasing in the agent's reciprocity so as to restore incentives to exert effort.

Next consider the case where the limited-liability constraint is binding. It is easy to verify that, in this case, the principal optimally sets the same bonus as in section 3.5 (corrected for the number of agents) and does not use promotion incentives. The reason is twofold. First, the problem that promotion incentives resolve in case the

[11] It should be taken into account that with two agents, first-best attention per worker is given by: $a = \frac{(H-L)\gamma}{2\theta\rho - \gamma^2}$.

agent's participation constraint is binding is non-existent when the limited-liability constraint binds, as the principal cannot recoup the agent's happiness gains from attention. Second, under limited-liability, promotion incentives are a more expensive instrument than bonus pay. The reason is that the principal always awards the promotion prize to one of the agents (even when both produce low output), which is costly when the limited-liability constraint binds.

3.7 Empirical analysis

This section empirically examines our theoretical predictions on the relationship between an employee's reciprocity and the type of incentive pay offered by his employer. Unfortunately, our data do not allow us to observe the strength of incentives workers receive. All that we know is *whether* workers receive a particular type of incentive. Therefore, we use our theoretical framework to derive predictions regarding an individual's likelihood of receiving a certain type of incentive scheme.

Our theory makes a clear prediction regarding promotion incentives: All else equal, workers who are more reciprocal should be more likely to receive promotion incentives. The reason is that promotion incentives do not dilute the principal's incentives to provide attention, and this advantage is particularly important for workers who are more reciprocal. We expect this relation to be particularly strong for workers in small firms, since then the manager is more likely to be residual claimant.[12]

Our theoretical predictions regarding individual bonus incentives are most clear for workers who do not receive promotion incentives. As we have seen in section 3.5, bonus pay dilutes the principal's incentive to provide attention, implying a negative relation between reciprocity and the likelihood of receiving bonus incentives. We expect this relation to be particularly strong in small firms, as managers in small firms are more likely to be residual claimant.

[12] If the manager is not a residual claimant, our results need not change qualitatively when the manager's income depends on his unit's profits. In large organizations, a lack of attention provision at the highest levels (from CEO to middle managers) may trickle-down to lower levels (from middle managers to workers). Shanock and Eisenberger (2006) find evidence for such effects.

We thus examine the following predictions:[13]

1. Workers who are more reciprocal are more likely to receive promotion incentives.
2. The relationship as described in prediction 1 is particularly strong for workers in small firms.
3. Among workers who do not receive promotion incentives, workers who are more reciprocal are less likely to receive bonus incentives.
4. The relationship as described in prediction 3 is particularly strong for workers in small firms.

For the empirical analysis we make use of the 2004 and 2005 waves of the German Socio-Economic Panel (GSOEP), a survey representative for the German population.[14] This dataset is unique in that it contains data on both worker's performance pay and worker's reciprocity. Reciprocity is measured by asking how well each of the three following statements applies: (1) If someone does me a favor, I am prepared to return it; (2) I go out of my way to help somebody who has been kind to me before; (3) I am ready to undergo personal costs to help somebody who helped me before. The extent of agreement with these statements is indicated on a 7-point scale, where 1 indicates profound disagreement and 7 means that the statement applies perfectly. Following Dohmen et al. (2009), we construct a measure of reciprocity by taking the average score on the three statements.[15] Our dependent variable (the worker's performance pay) is measured by asking whether people's job performance is regularly assessed by a supervisor and whether this performance evaluation has consequences

[13] As we have seen in the previous section, our theory also predicts that when a limited-liability constraint is binding, the relation between reciprocity and the likelihood of receiving promotion incentives will be weaker. We analyze this in the working paper version of this chapter (see Dur et al. 2008). Using union-membership and low income as proxies for limited-liability, we find evidence in line with these predictions. However, we should be careful with interpreting these results, because the proxies for limited liability may also pick up other differences in unobserved characteristics.

[14] Detailed information about the GSOEP can be found at www.diw.de/gsoep/. By merging the 2004 and 2005 waves, we lose some observations. For this reason, the two last rows in table 1 do not add up to the totals in the first row. For the same reason, we lose some additional observations in the regressions.

[15] Dohmen et al. (2009) distinguish between positive reciprocity and negative reciprocity, where the former refers to the inclination to reciprocate kind acts and the latter to reciprocation of unkind acts. Because we consider attention as a gift to the agent, we focus on positive reciprocity. For further discussion of these measures of reciprocity, see Dohmen et al. (2009).

for promotion and/or for receiving a yearly bonus.[16] A value of 1 indicates a positive answer. In our analysis, we leave out those individuals who indicate that they do not get a regular performance evaluation by their supervisor. The reason is that some of these individuals, although lacking a formal performance evaluation, in fact may receive bonus or promotion incentives, which we cannot observe in the data.[17] For example, we cannot observe whether an individual receives bonus incentives based on objective criteria. We provide some summary statistics of the two key variables in figure 1 and table 1, the control variables are described in table 2.

We examine the first prediction by estimating two Probit-equations, as shown in the first two columns of table 3.[18] The difference between the two equations is that in the second column, we do not control for firm size, industry, and occupation. In line with prediction 1, the coefficient for reciprocity is positive and statistically significant at the 1% level. Also, the effect of reciprocity is economically significant: a 1-point increase on the scale of reciprocity increases the probability of getting promotion incentives by more than 5 percentage points on average (the mean marginal effect follows from multiplying the coefficient with the appropriate reduction factor). Taking into account the fact that about 45% of the sample used in the regression gets promotion incentives, this is quite a large effect. The effect remains highly significant but reduces somewhat in size when we do not control for firm size, industry, and occupation. This suggests that if sorting mechanisms are present, they are rather subtle: reciprocal workers do not sort on the basis of occupation, industry or

[16]Unfortunately, the survey question does not rule out that a bonus is awarded based on relative performance evaluation instead of an absolute criterion. This may lead to an upward bias in the estimates of the relation between reciprocity and the use of bonus pay.

[17]The main conclusions are qualitatively robust to inclusion of individuals who do not get a performance evaluation. However, the picture often becomes blurred in the sense that the t-values of the coefficient for the reciprocity variable fall, which is not surprising if, indeed, some workers who do not get a performance evaluation do receive performance pay. We also checked whether our results are sensitive to inclusion of other personality characteristics (the Big 5 and locus of control), additional regional dummies, hours of work, and wage income. Including these variables does not change our results. Correcting for wage income does not lead to significant drops in the estimates in table 3 and 4, which might have been expected as reciprocity is positively correlated with wage income, see Dohmen et al. (2009). The reason for not including wage income and hours worked in the reported regressions is the endogeneity of both variables. Also, limiting the sample to white-collar workers does not significantly affect our estimates.

[18]Except for the inclusion of the reciprocity parameter and the distinction between promotion incentives and bonus pay, the specification is similar to Grund and Sliwka (2010).

firm size.

We examine the second prediction by re-estimating the equations for a sample of small and large firms. A firm is considered ‘small’ if it has less than 100 employees and ‘large’ if it has 100 employees or more. The results are shown in columns (3) and (4) of table 3. Clearly, reciprocity has a larger effect on the probability of receiving promotion incentives in small firms than in large firms, which is supportive of prediction 2. However, the difference between the coefficients is not statistically significant. A similar picture arises when we leave out some of the controls and when we replicate the regression for firms with less than 20 employees (approximately 260 observations).

The third prediction is that among workers who do not receive promotion incentives, workers who are more reciprocal are less likely to receive bonus incentives. In the first two columns of table 4, we regress the likelihood of receiving bonus pay on reciprocity for these workers. Controlling for firm size, industry, and occupation, the coefficient for reciprocity has the predicted sign but is highly insignificant (see the first column). The significance does not improve if we drop the controls, as can be seen from the second column of table 4. Worse still, the coefficient now also has the wrong sign. We learn two things from these regressions. First, just as in the regression of promotion incentives on reciprocity, sorting seems to play no role. Second, we find no evidence in favor of prediction 3. Next, we split the sample into small and large firms, where 100 employees is again taken as the cut-off point. The third and fourth column of table 4 show the results. The reciprocity coefficients are insignificant and the signs are opposite to prediction 4. These findings are robust to dropping (sets of) controls, considering firms with less than 20 employees, as well as to adding workers who do receive promotion incentives to the sample.

3.8 Concluding remarks

We have analyzed how incentives for a worker and his superior interact using a gift-exchange model where the worker is reciprocal to attention of his superior. In our model, attention is the superior’s input in the exchange relation, whereas the worker reciprocates with higher effort. This reciprocity directly links the inputs of a manager to that of his worker, so that production becomes a joint effort. We have studied a common trade-off in models of double moral-hazard, where stronger incentives for one of the partners weaken those for the other. The central result is that bonus pay for the worker weakens the incentives for his superior to provide attention. We have seen that promotion incentives are particularly helpful when workers are reciprocal, since it takes away the commitment problem for the superior in providing attention. This last result is supported by data from the German Socio-Economic Panel, which show a clear positive relation between a worker's reciprocity and the likelihood of receiving promotion incentives, even after controlling for a rich set of observable characteristics.

Obviously, in practice, several other aspects of organizational life need to be taken into consideration when choosing the type of incentives offered to workers. While our analysis has focused on the quality of relationships between managers and workers, the quality of co-worker relationships may also be very important for worker's job satisfaction and, hence, for an organization's ease in attracting and retaining workers. Promotion incentives may hurt or help in this respect, see among others Grund and Sliwka (2005) and Dur and Sol (2010).

3.9 Tables and figures

Figure 1: Frequency distribution of reciprocity for employees who get a regular performance appraisal

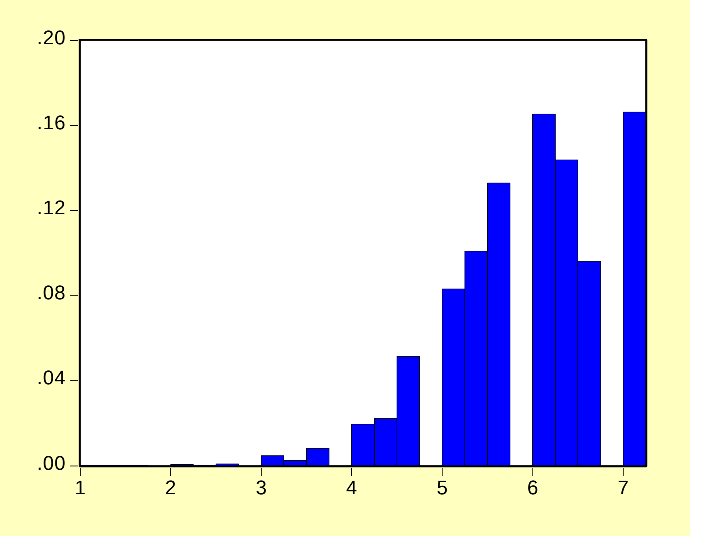

Note: the frequency distribution for all employees (including those not receiving a regular performance appraisal) is very similar to the one shown above. Details are available upon request.

Table 1: Summary statistics on performance pay

Subgroup		Frequency			Subtotal
No performance appraisal		72% (8179)			
Performance appraisal		28% (3159)			
Appraisal has consequences for:	Bonus only	Promotion only	Bonus and Promotion	Neither	
Total	14% (452)	25% (803)	19% (586)	42% (1318)	100% (3159)
Small firms (<100 employees)	14% (102)	17% (123)	12% (84)	57% (405)	100% (714)
Large firms (≥100 employees)	14% (316)	29% (634)	22% (481)	35% (769)	100% (2200)

Table 2: Description of independent variables used in regression

Risk Attitude	Willingness to take risks on a scale $0-10$, where 0 is extremely risk averse.
Female	Dummy variable: 1=female.
Age	Age in years.
Years of education	Years of education.
Tenure	Years employed by the firm.
Part-time	Dummy variable: 1=part-time job.
East-Germany	Dummy variable: 1 if the place of work is in East-Germany or East-Berlin.
Firm size dummies	Number of employees n in whole company, categorized as follows: (1) $n < 5$ (basis) (2) $5 \leq n < 20$, (3) $20 \leq n < 100$, (4) $100 \leq n < 200$, (5) $200 \leq n < 2000$, (6) $n \geq 2000$.
Industry dummies	One-digit industry code: 1=Agriculture, 2=Energy, 3=Mining, 4=Manufacturing, 5=Construction, 6=Trade, 7=Transport, 8=Bank and Insurance, 9=Services, 10=Other (basis).
Occupational dummies	The individual's occupation and occupational level: Blue collar worker (5 levels), white collar worker (6 levels) or civil servant (4 levels).

Table 3: Effect of reciprocity on the probability of receiving promotion incentives. (Robust standard errors in parentheses)

Dependent variable:	Promotion incentives		Promotion incentives	
			Firm size <100	Firm size ≥100
	(1)	(2)	(3)	(4)
Reciprocity	0.154***	0.100***	0.227***	0.139***
	(0.032)	(0.029)	(0.072)	(0.037)
Risk attitude	-0.013	-0.005	-0.045*	-0.002
	(0.013)	(0.012)	(0.027)	(0.016)
Female	-0.117*	-0.197***	-0.190	-0.105
	(0.068)	(0.057)	(0.152)	(0.076)
Age	0.057***	0.093***	0.048	0.063**
	(0.021)	(0.015)	(0.041)	(0.025)
Age squared	-0.001***	-0.001***	-0.001*	-0.001***
	(0.000)	(0.000)	(0.001)	(0.000)
Years of education	0.005	0.111***	-0.026	0.012
	(0.013)	(0.009)	(0.032)	(0.015)
Tenure	0.002	0.025***	-0.013	0.004
	(0.004)	(0.003)	(0.009)	(0.004)
Part-time	-0.178*	-0.208**	-0.106	-0.175
	(0.092)	(0.081)	(0.189)	(0.107)
East-Germany	-0.392***	-0.533***	-0.505***	-0.364***
	(0.072)	(0.065)	(0.143)	(0.085)
Dummies for Firm size, Industry, and occupation	Yes	No	Yes	Yes
Observations	2726	2829	631	2088
Pseudo R^2	0.228	0.097	0.246	0.207
Reduction factor	0.369	0.362	0.372	0.372

Constant is included but not reported.

***,**,* indicate significance at 1%, 5%, and 10% level respectively.

Table 4: Effect of reciprocity on the probability of receiving bonus pay. (Robust standard errors in parentheses)

Dependent variable:	Bonus pay		Bonus pay Firm size <100	Bonus pay Firm size ≥100
	(1)	(2)	(3)	(4)
Reciprocity	-0.011	0.001	0.102	-0.056
	(0.044)	(0.042)	(0.092)	(0.053)
Risk attitude	0.015	0.021	-0.031	0.003
	(0.019)	(0.018)	(0.035)	(0.023)
Female	-0.193**	-0.362***	-0.226	-0.159
	(0.094)	(0.082)	(0.186)	(0.111)
Age	0.074***	0.110***	0.097**	0.058*
	(0.027)	(0.02)	(0.048)	(0.034)
Age squared	-0.001**	-0.001***	-0.001**	-0.001
	(0.000)	(0.000)	(0.001)	(0.000)
Years of education	0.043**	0.058***	0.060	0.029
	(0.019)	(0.013)	(0.041)	(0.023)
Tenure	0.007	0.008*	0.013	0.003
	(0.005)	(0.004)	(0.010)	(0.006)
Part-time	-0.182	-0.345***	0.001	-0.276*
	(0.127)	(0.119)	(0.231)	(0.156)
East-Germany	-0.094	-0.176**	-0.216	-0.017
	(0.094)	(0.086)	(0.161)	(0.120)
Dummies for Firm size, Industry and occupation	Yes	No	Yes	Yes
Observations	1461	1474	448	1013
Pseudo R^2	0.160	0.078	0.196	0.165
Reduction factor	0.3766	0.3785	0.3698	0.3883

Constant is included but not reported.

***,**,* indicate significance at the 1%, 5%, and 10% level respectively.

Chapter 4

Dynamic Incentive Effects of Relative Performance Pay

Joint with Josse Delfgaauw, Robert Dur, and Willem Verbeke.

4.1 Introduction

Non-linear pay-for-performance plans have dynamic incentive effects when employees receive intermediate performance information over the course of the incentive period. For instance, consider a salesman who can earn a bonus by attaining a monthly sales target while receiving daily or weekly sales figures. When realised sales during the month are such that it remains challenging but possible to reach the target, the bonus scheme provides strong incentives. The incentive effect is much weaker, however, when realised sales during the month are particularly high or low. High intermediate sales imply that the salesman can hardly miss the bonus, while low intermediate sales imply that the target is practically out of reach.

More generally, workers can use intermediate performance information to determine how much additional performance is necessary to obtain a bonus. This creates dynamic incentive effects, where the incentive effect of the pay-for-performance plan at each point in time depends on realised performance until then. Incentive plans based on relative performance, where prizes are awarded for outperforming sufficiently many competitors, are particularly prone to dynamic incentive effects. Sports leagues are a common example. In the workplace, examples range from employee-of-the-month contests, to beat-the-index bonuses for stock brokers, and to job promotion contests.[1]

Casas-Arce and Martinez-Jerez (2009) show formally that for contests with a large number of participants, the incentive effect of a relative performance incentive scheme is hump-shaped in lagged relative performance. Competitors who find themselves trailing far behind may perceive catching up to be impossible and consequently give up trying. Similarly, competitors who are far ahead may perceive losing as impossible and slack off as well. In contrast, incentives are highly salient for competitors who find themselves almost tied in intermediate performance. Analyzing sales contests among retailers of a commodities company, Casas-Arce and Martinez-Jerez (2009) find indeed that competitors in winning positions reduce performance when their lead increases. However, the performance of trailing competitors does not decrease when they lag further behind. The authors conjecture that this result might be affected by attrition bias. Frank and Obloj (2011) do find the predicted hump-shaped pattern in their analysis of a competition among units of a retail bank.[2]

Testing for the presence and strength of dynamic incentive effects is hampered by two issues. First, in contests with a limited number of participants, a competitor's optimal strategy depends on (its perception of) its competitors' strategies. A trailing competitor may be best off by accepting its loss when the other competitors keep effort high, but not when they would slack off. Second, serial correlation in performance biases estimates of the effect of intermediate relative performance on subsequent performance. For instance, positive serial correlation would imply that a

[1] As a concrete example, more than half of the remuneration of the executive directors of oil company Shell is based on a ranking of Shell's performance relative to its four main competitors on four publicly available measures. The incentive plan has a three-year horizon, during which the companies regularly release the latest figures with respect to these performance measures (Royal Dutch Shell, 2009).

[2] Relatedly, Fershtman and Gneezy (2009) let kids run side-by-side and find that increasing incentives yield higher performance but also a higher fraction of kids giving up during the race. Following the early literature on tournament theory (Lazear and Rosen, 1981, Green and Stokey, 1983, Nalebuff and Stiglitz, 1983), most of the literature has abstracted from dynamic incentive effects of tournaments. A recent string of theoretical papers studies the cost and benefit to a principal of providing intermediate relative performance feedback during a contest between his agents (Aoyagi 2010, Ederer 2010, Gershkov and Perry 2009, Goltsman and Mukherjee 2010).

positive shock to performance in the previous period increases both relative intermediate performance and current performance. Casas-Arce and Martinez-Jerez (2009) and Frank and Obloj (2011) employ a method developed by Arellano and Bond (1991) that relies on taking first differences and using lagged values of independent variables as instruments of the independent variables.

We take a unique approach in tackling both issues by setting up a relative performance pay scheme where only one of the 'competitors' can earn a prize, while the other participants are kept unaware of their involvement. This implies that the strategies of all non-competing participants are exogenous, allowing us to use their performance as an instrument for intermediate relative performance of the competing participant. More specifically, we study the dynamic incentive effects of this relative performance pay scheme by conducting a natural field experiment in a Dutch retail chain. We provide the employees of 93 stores randomly selected from 189 of the company's stores with the opportunity to earn a bonus. The bonus is awarded when a treatment store outperforms three comparable stores *from the control condition* over the course of a four-week period (February 2010). Each week, treatment stores receive a poster with the performance of all four stores in their group. Importantly, the employees of the three comparison stores cannot earn a bonus, do not learn that another store can earn a bonus by beating their performance, and do not receive any relative performance feedback. This way the treatment stores compete against stores that are not competing. This allows us to use the performance of the three comparison stores as an instrument for trailing behind or being ahead: lagged performance of the comparison stores does affect intermediate relative performance, but does not affect current performance of the treatment store other than through lagged relative performance. Hence, using this instrument, our estimates are not biased by serial correlation in stores' own performance.

Our results are as follows. First, we find a positive effect of intermediate relative performance on current performance for stores close to the target, particularly in the last two weeks of the experiment. This effect is substantial: a one percentage point increase in intermediate relative performance increases current performance by 0.72 percent. Stores lagging far behind do not respond to intermediate relative performance. This suggests that the employees in these stores gave up trying to win. Hence, as predicted by theory, we find that intermediate relative performance matters more for competitors that perform close to target than for competitors that lag far behind. During the contest, hardly any treatment store managed to get far ahead of all its comparison stores. Hence, we cannot test the hypothesis that high-performers slack off as their lead increases.

Second, we find no average treatment effect of introducing the contest, neither for the four weeks taken together nor for one of the weeks separately. This contrasts with several recent findings on the incentive effects of tournaments. In another retail chain, we do find a substantial positive effect of introducing a standard tournament among shops (Delfgaauw et al., 2009), as do Erev et al. (1993) and Bandiera et al. (2009) among teams of fruit pickers and Casas-Arce and Martinez-Jerez (2009) among retailers of a commodities company. Even more striking, several recent papers suggest that the mere provision of relative performance feedback can be sufficient to trigger higher performance (Azmat and Iriberri (2010), Blanes i Vidal and Nossol (2009), Delfgaauw et al. (2009), Kosfeld and Neckermann (2010)). Bandiera et al. (2009) obtain an opposite result. A possible explanation for our divergent result is that beating unaware contestants, as in our setting, is less exciting than beating competing contestants.

Our experiment involves one incentive period of four weeks. When incentive schemes are repeated over time, as with monthly or year-on-year targets, other types of dynamic incentive effects may arise. For instance, sales may be shifted forward or backward in time around the incentive commencement date in order to meet the current target or to alleviate the difficulty of meeting the next target; see Asch (1990) and Oyer (1998) for empirical evidence and Cadsby et al. (2010) for a related lab experiment. Furthermore, when the targets in repeated incentive schemes are based on historical performance, workers have an incentive to beat the target by only a limited amount even it would be possible to greatly outperform the target. Bouwens and Knoops (2010) find evidence in line with such ratchet effects, using store-level data from a retail chain. Cooper et al. (1999) and Charness et al. (2010) find ratchet effects in the lab. Ratchet effect considerations may explain why we find no average treatment effect, as workers may have feared that a strong response to the introduction of the relative performance pay scheme would result in higher targets in their regular incentive scheme.

4.2 Experimental design

The experiment took place in February 2010 in a retail chain in the Netherlands that sells computer games, music, and movies. At the start of 2010, the retail chain owned 208 geographically dispersed stores, operating under two different brands. Each store employs on average 5 employees, including a store manager. The company's central management decides on the range of products sold, pricing, and advertisement. Store managers are responsible for day-to-day operations. Employees receive rather weak incentive pay on top of their base salary, based on their shop's yearly sales growth and a subjective performance evaluation. The company's management was not satisfied with the effects of this incentive scheme and wished to learn more about the effects of short-term incentives, in particular of sales contests. The pre-existing incentive scheme remained in place during the experiment.

We designed a relative performance incentive scheme to be implemented in a randomly selected subset of stores (the treatment condition), while the rest of the stores comprised the control condition. All employees (including the shop manager) of a store in the treatment condition could earn a bonus by sufficiently outperforming three preselected stores from the control condition. Stores in the control condition could not earn a bonus, and employees in the treatment stores were informed about this. Performance is measured as cumulative sales revenue in percentage deviation of budgeted sales in February 2010 (a period of 4 weeks).[3] Let $y_{s,w}$ be sales and $b_{s,w}$ budgeted sales of store s in week w, respectively. Weekly performance $p_{s,w}$ is given

[3] The budgeted sales are forecasts for shops' weekly sales as determined by the company's management in October 2009 (at the start of the financial year) for a year onwards. These budgeted sales boil down to a forecast for total sales of the whole chain, with each store expected to bring in a fixed share of total sales. Hence, a combination of week and store fixed effects explains all variation in the log of budgeted sales in our data. The company gives shop managers weekly feedback on sales relative to budgeted sales, which makes it a natural measure of performance.

by

$$p_{s,w} = \frac{y_{s,w} - b_{s,w}}{b_{s,w}} \cdot 100\% \tag{4.1}$$

and cumulative performance over February 2010 is given by

$$p_s^{CU} = \frac{\sum_{w=E1}^{E4} y_{s,w} - \sum_{w=E1}^{E4} b_{s,w}}{\sum_{w=E1}^{E4} b_{s,w}} \cdot 100\% \tag{4.2}$$

where the summation is over the four experimental weeks $E1$ to $E4$ (i.e. week 5, 2010 to week 8, 2010).

All employees of a treatment store received a bonus of gross 150 euro when their shop's performance in February 2010 was at least 10 percentage points higher than the performance of all three comparison stores. When a treatment store scored between 5 and 10 percentage points above all three comparison stores, its employees received 75 euro.[4] Lastly, outperforming all three comparison stores by less than 5 percentage points yielded a cake for the treatment store, but only if the treatment store also performed above budget.[5]

All communication on the experiment towards the shops went through the company's regular channels, so shop managers and employees were not aware of our involvement. Hence, our experiment classifies as a natural field experiment (Harrison and List, 2004). In January 2010, the company informed all stores that a randomly selected set of stores would get the opportunity to earn a bonus in February 2010, and that all other stores could look forward to a similar opportunity later that year. On January 22, the treatment stores were informed about the details of the relative performance incentive scheme.

During the experiment, we provided weekly feedback to the treatment stores on their relative performance in the form of a poster. The poster contained the cumulative sales relative to budget figures of the treatment shop and its three comparison shops, ranked in descending order. Furthermore, on Monday February 1, all treat-

---

[4] For employees who did not have a full-time contract, the size of the bonus was proportional to the contractual number of hours. Hourly wages for personnel in the shops are close to the minimum wage, which makes that receiving the high bonus would increase monthly earnings by about 10%.

[5] The latter requirement only applied for the cake, not for any of the two bonusses. This requirement was a last-minute addition by the company's management to the rules.

ment stores received a large poster, with room to glue on the four posters with weekly rankings to be received in the following weeks. Store managers were instructed to put up these posters in the store's canteen.[6] Stores in the control condition did not receive posters, nor any other type of relative performance information.

Our design has two advantages over a regular competition. First, as treatment stores only receive a bonus when they outperform comparable stores from the control condition, the payout is relatively low when the incentive has little effect on performance. This was seen as a major benefit by the company's management. Second, performance of the comparison stores is exogenous to the incentive scheme, as these stores could neither earn a bonus, nor received any relative performance feedback, and were not aware that their performance played a role in the incentive scheme. We exploit differences in comparison stores' performance during the experiment to analyse how treatment stores' intermediate relative performance affects the effect of the incentive scheme in subsequent weeks.

We used weekly sales and budget data of 194 stores[7] for the weeks 40 to 53 in 2009 to assign stores to the treatment and control conditions, and to match treatment stores with comparable stores from the control condition, as follows. First, we created four equally large strata based on store size as measured by average weekly sales revenues. Randomly, half of the stores in each stratum was assigned to the treatment condition, while the remaining half of the stores were assigned to the control condition. Subsequently, we matched each treatment store to three control stores from the same stratum. To reduce opportunities for collusion, we imposed that each treatment store was matched to control stores located in other regions, as there is frequent communication between stores within a region. Apart from this regional separation, treatment stores were matched to the control stores that were most comparable in terms of the performance measure (cumulative sales revenue relative to the budget) for the period of week 40 to week 53 in 2009. Note that a control store can be matched to multiple treatment stores. After this assignment

[6] The company's regional managers were instructed to verify that all store managers actually put up the posters in the canteen. We have not heard about a single store manager who refused to do so.

[7] The company's management excluded a specific group of 14 stores from participating in the experiment.

procedure, we excluded one treatment store from the experiment as its budget figures turned out to be too unreliable. Furthermore, 3 treatment stores and 1 control store were shut down in January 2010. This leaves us with 93 stores in the treatment condition and 96 stores in the control condition. For each of these stores, we have weekly sales and budget figures for a period of in total 22 weeks, from week 40 in 2009 to week 8 in 2010. In some estimations, we separate the stores by size, by collapsing the two strata with the biggest stores together as well as the two strata with smallest stores. This yields 97 large stores and 92 small stores. Figures 1 and 2 show weekly sales and weekly performance, respectively, averaged over all stores. Average weekly sales show two spikes in December 2009, related to Sinterklaas and Christmas festivities, respectively. Removing these weeks from the analysis does not affect the results. Average performance hovers between plus and minus 20 percent. The spikes in sales are anticipated by the company's management when determining budgeted sales, as the spikes in sales do not carry over to average performance.

The descriptive statistics in Table 1 show that average sales does not differ between treatment stores and control stores, neither for the whole period nor for the first 14 weeks in the data used to stratify the stores. The same holds for budgeted sales and for performance as measured by (4.1). Note that on average, sales are below budget, but that variation in average performance across stores is large. Further, the number of employees per store does not differ significantly between the treatment and control stores. Lastly, in week 7 of 2010, a total of 29 stores were closed for one or two days in relation to carnival festivities, mainly in the south of the Netherlands. Treatment stores were slightly more often closed than control stores, but not significantly so. In all estimations below, we correct for the effect of carnival.

As a first hint of the overall effect of the relative performance incentive, Table 1 shows that there is no difference in average sales figures between treatment and control stores for the weeks with the bonus scheme (week 5, 2010 to week 8, 2010). Figures 3 and 4 provide further insight into the overall treatment effect, by plotting the differences between the treatment and control condition in average sales and in average performance, respectively, by week. The experiment took place in the final

Table 4.1: Descriptive statistics

	All stores		Treatment stores		Control stores	
	mean	Std	Mean	Std	Mean	Std
Sales	100.00	40.31	100.04	43.77	99.96	36.89
Sales weeks 40/2009 - 53/2009	115.18	46.62	115.27	50.81	115.08	42.43
Sales weeks 5/2010 - 8/2010	71.67	29.90	71.56	31.85	71.78	28.05
Budgeted sales	104.56	41.54	105.33	45.42	103.81	37.63
Budgeted sales weeks 40/2009 - 53/2009	122.90	48.82	123.80	53.38	122.02	44.23
Budgeted sales weeks 5/2010 - 8/2010	78.52	31.19	79.10	34.10	77.96	28.86
Performance	-0.03	0.12	-0.04	0.11	-0.02	0.12
Performance weeks 40/2009 - 53/2009	-0.06	0.11	-0.07	0.10	-0.05	0.12
Performance weeks 5/2010 - 8/2010	-0.08	0.14	-0.09	0.13	-0.08	0.14
Number of employees	5.45	1.99	5.24	1.69	5.66	2.22
Number of days closed for carnival (week 7)	0.25	0.63	0.19	0.56	0.31	0.69
Number of stores	189		93		96	

Notes:
Performance is defined as (sales-budgeted sales)/budgeted sales.
For confidentiality reasons, sales and budgeted sales figures are indexed to the average sales per store per week over the whole sample.
None of the differences between treatment stores and control stores are significant at the 10%-level.

Figure 4.1: Average sales per store

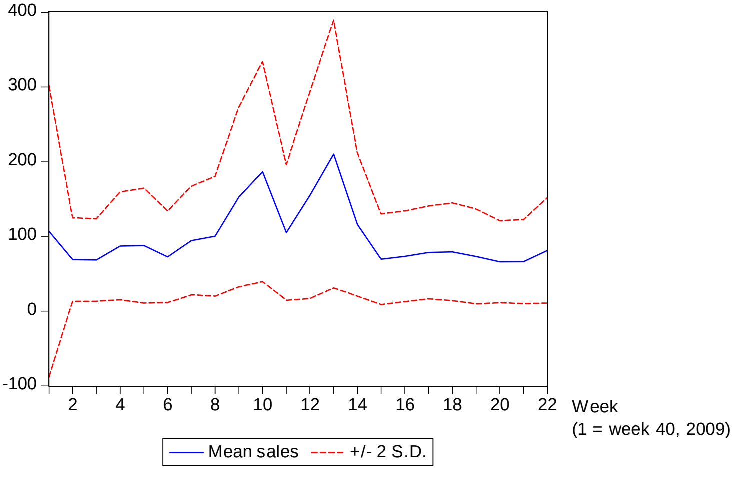

Sales is indexed, where average sales per store-week over the whole period = 100.

Figure 4.2: Average weekly performance

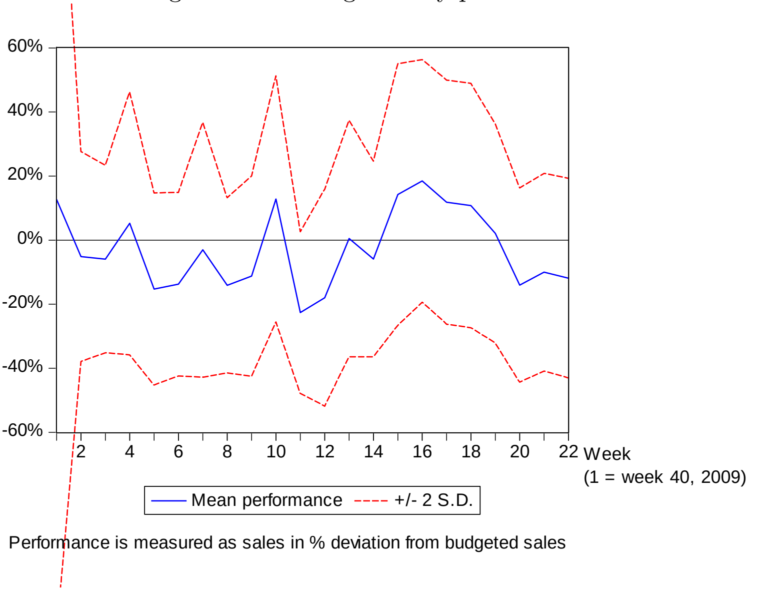

Performance is measured as sales in % deviation from budgeted sales four weeks of the period shown. Both figures show no sign of a positive treatment effect, possibly with the exception of the final week. A second hint of the overall effect is given by the fact that only 13 stores earned a prize: 5 stores earned the high bonus, another 5 stores earned the low bonus, and three stores were entitled to cake.

4.3 Method

We assess the average effect of the relative performance incentive scheme using OLS with week and store-fixed effects, by estimating

$$\ln(y_{s,w}) = \alpha_s + \theta_w + \gamma B_{s,w} + \kappa F_{s,w} + \varepsilon_{s,w} \tag{4.3}$$

where $\ln(y_{s,w})$ is the log of sales of store s in week w. Store and week-fixed effects are given by α_s and θ_w, respectively. $B_{s,w}$ is a dummy variable that takes the value one for stores in the treatment condition from week 5 to week 8 in 2010, $F_{s,w}$ measures the number of days shop s is closed for carnival festivities in week w (this variable takes positive values only in week 7, 2010), and $\varepsilon_{s,w}$ is an error term. We estimate the effect on sales rather than on the performance measure (4.1) used in the incentive scheme. Budgeted sales are set in advance by the company's central management and cannot be affected by stores. This implies that shops can affect their performance only through sales. As effects on sales are more easily interpreted, we use sales to measure the effects of the relative performance incentive scheme.[8]

[8]Using performance (4.1) as dependent variable instead yields similar results. This is unsurprising given the way stores' budgeted sales are determined, see footnote 3.

Figure 4.3: Average sales of treatment stores divided by average sales of control stores

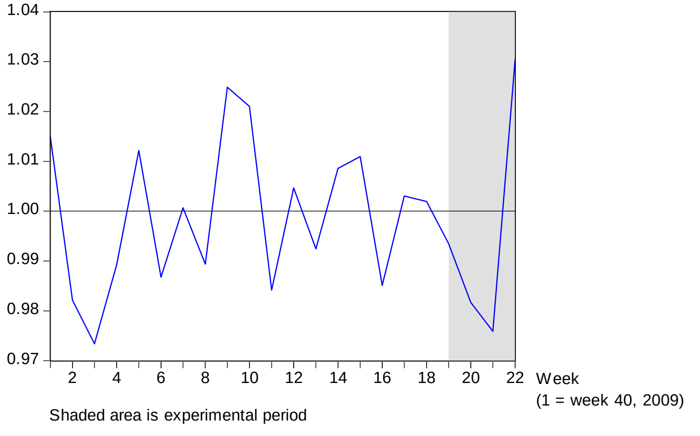

The main goal of this study is to analyse the effect of intermediate relative performance on subsequent performance. First, we introduce some additional notation. Let T and C be the sets of stores in the treatment and the control condition, respectively. Further, denote by $c_t \in C$ a control store matched to treatment store $t \in T$. Lastly, let $p^{CU}_{s,w-1}$ denote the *cumulative* performance of store s during the experiment up to but not including week w, as measured by cumulative sales over budget in February 2010:

$$p^{CU}_{s,w-1} = \frac{\sum_{w=E1}^{w-1} y_{s,w} - \sum_{w=E1}^{w-1} b_{s,w}}{\sum_{w=E1}^{w-1} b_{s,w}} \cdot 100\%. \tag{4.4}$$

Hence, $p^{CU}_{s,w-1}$ is the performance figure for store s as depicted on the poster received at the start of week w during the experiment. The effect of intermediate performance of treatment stores relative to the best-performing comparison store on subsequent sales can be estimated by

$$\ln(y_{s,w}) = \alpha_s + \theta_w + \gamma B_{s,w} + \mu \left(p^{CU}_{t,w-1} - \max_{c_t} \left[p^{CU}_{c_t,w-1} \right] \right) B_{s,w} + \kappa F_{s,w} + \varepsilon_{s,w} \tag{4.5}$$

Figure 4.4: Difference in average performance between treatment stores and control stores

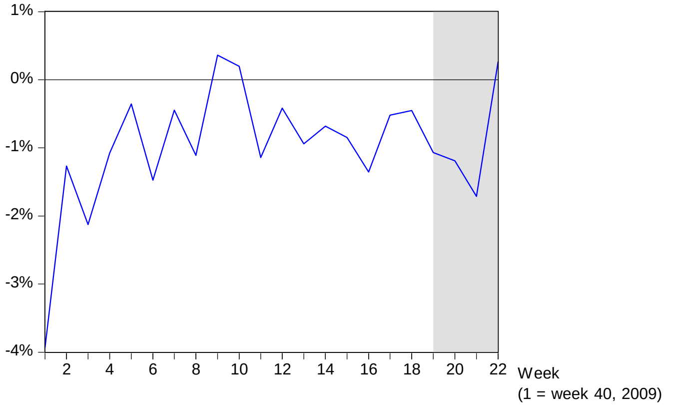

Performance is measured as sales in % deviation from budgeted sales
Shaded area is experimental period

where the term $p^{CU}_{t,w-1} - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$ gives the difference in cumulative performance during the experiment between treatment store t and its best-performing comparison store c_t up to and including the previous week.[9] Since the experiment lasted four weeks, we have three intermediate relative performance figures per treatment store, corresponding to a total of 279 treatment store-week observations. Control stores cannot earn a bonus and do not receive posters with rankings. Hence, $\max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$ is fully exogenous to $y_{t,w}$. However, in case of serial correlation in the error structure of sales, $p^{CU}_{t,w-1}$ is correlated with $y_{s,w}$ (see (4.4)). Estimating (4.5) without taking account of serial correlation would yield a biased estimate of μ. Therefore, we instrument the difference in intermediate performance $p^{CU}_{t,w-1} - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$ by the expected difference

$$D_{t,w-1} = E\left[p^{CU}_{t,w-1}\right] - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]. \tag{4.6}$$

The expected cumulative performance of treatment store t in the experiment $E\left[p^{CU}_{t,w-1}\right]$ is set equal to the average performance in the 18 weeks prior to the start of the ex-

[9] This variable is set to zero for control stores.

periment (week 40, 2009 to week 4, 2010) while accounting for week-fixed effects in performance during the experiment:

$$E\left[p_{t,w-1}^{CU}\right] = \frac{1}{18}\sum_{w=1}^{18} p_{t,w} + \frac{\sum_{w=E1}^{w-1} b_{t,w}\theta_w^p}{\sum_{w=E1}^{w-1} b_{t,w}} \tag{4.7}$$

where θ_w^p is the week-fixed effect from estimating

$$p_{s,w} = \alpha_s^p + \theta_w^p + \gamma^p B_{s,w} + \kappa^p F_{s,w} + \varepsilon_{s,w}^p$$

with superscript p denoting that the estimates relate to performance as dependent variable.[10] Most importantly, this implies that for each treatment store, variation in $D_{t,w-1}$ across experimental weeks stems solely from variation in $p_{ct,w-1}^{CU}$ and in the weighted week-fixed effects, which are both unrelated to $\varepsilon_{t,w}$ given the design of our experiment.

Equation (4.5) estimates a linear effect of intermediate relative performance. However, the incentive scheme is likely to have the biggest effect when treatment stores learn that they are close to the relative performance targets for winning a bonus (Casas-Arce and Martinez-Jerez, 2009).[11] Treatment stores lagging far behind in the intermediate ranking may give up, and treatment stores far ahead may reduce their efforts when they anticipate that they can hardly miss the bonus. In the course of the experiment, we have many treatment stores that face an uphill battle, while there are only few stores that are comfortably ahead. In total, we have only 8 store-week observations where treatment stores' intermediate relative performance is more than 10 percentage points above the target for the high bonus (i.e. with $p_{t,w-1}^{CU} - \max_{ct}\left[p_{ct,w-1}^{CU}\right] > 0.20$). This implies that we cannot test whether stores that greatly outperform their comparison stores reduce their efforts.[12] We can test whether stores that lag far behind reduce their efforts, by allowing the effect of

[10] We weight the week-fixed effects by budgeted sales $b_{t,w}$ to account for the fact that weeks with a higher absolute budgeted sales volume have a higher weight in cumulative performance, see (4.4).

[11] The theory developed by Casas-Arce and Martinez-Jerez (2009) predicts that performance is hump-shaped in intermediate relative performance, but does not predict the exact level of relative performance at which the incentive effect peaks.

[12] Excluding these 8 observations from the analysis does not affect any of the results.

intermediate relative performance on current performance to differ between stores that lag far behind and stores that are still in the running.

In determining which stores still have a chance of earning a bonus, we cannot use the actual difference between the lagged cumulative performance of the treatment store and its best control, as given by (4.4). Serial correlation in $y_{t,w}$ would bias the estimates. Hence, we again use the estimated difference (4.6) to determine stores' chances of earning a bonus. Rather arbitrarily, we set the bar for being too far behind at a 5 percentage point lag relative to the best performing comparison store. Note that stores that lag 5 percentage points behind need to improve their relative performance by 5 percentage points in order to win a cake and by at least 10 percentage points to obtain a bonus. We do vary the bar to assess the robustness of the results. Let $I_{t,w-1}$ be a dummy that takes value 1 for treatment stores whenever $D_{t,w-1} > -0.05$ and zero otherwise. This yields 52 store-week observations where $I_{t,w-1} = 1$, out of a total of 279 treatment store-week observations with intermediate relative performance figures. We estimate

$$\begin{aligned} \ln(y_{s,w}) &= \alpha_s + \theta_w + \gamma B_{s,w} + \mu \left(p_{t,w-1}^{CU} - \max_{c_t} \left[p_{c_t,w-1}^{CU} \right] \right) B_{s,w} + \\ &\quad + \delta I_{t,w-1} B_{s,w} + \nu \left(p_{t,w-1}^{CU} - \max_{c_t} \left[p_{c_t,w-1}^{CU} \right] \right) I_{t,w-1} B_{s,w} + \kappa F_{s,w} + \varepsilon_{s,w} \end{aligned} \tag{4.8}$$

again instrumenting the difference in intermediate performance $p_{t,w-1}^{CU} - \max_{c_t} \left[p_{c_t,w-1}^{CU} \right]$ by the expected difference $D_{t,w-1}$ as given by (4.6).

In all of our estimations, we cluster standard errors at the store level to correct for serial correlation within stores, as well as for heteroscedasticity across stores (see Bertrand et al. (2004) for a discussion of the importance of correcting for serial correlation in Difference-in-Difference estimation).

4.4 Results

The first column in Table 2 gives the results of estimating (4.3). On average, the relative performance incentive scheme did not affect sales. The second column of Table 2 shows that there is some variation in the estimated treatment effect by week,

Table 4.2: Average treatment effect

	Dependent variable: ln(sales)	
	(1)	(2)
Treatment	-0.004	
	(0.013)	
Treatment week 1		-0.003
		(0.019)
Treatment week 2		-0.013
		(0.020)
Treatment week 3		-0.020
		(0.017)
Treatment week 4		0.021
		(0.019)
Carnival	-0.026*	-0.028*
	(0.015)	(0.015)
Store-fixed effects	yes	yes
Week-fixed effects	yes	yes
Store-week observations	4158	4158
Stores	189	189
R^2	0.9281	0.9281

Standard errors clustered at the store level in parentheses.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

but none of the estimates differs significantly from zero.

The first three columns in Table 3 present the results of estimating (4.5) using OLS and IV-2SLS, respectively, with $D_{t,w-1}$ instrumenting $p^{CU}_{t,w-1} - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$. The second column contains the first-stage regression of the IV-estimation. Actual intermediate relative performance increases one-for-one with our instrument, predicted difference (4.6). This instrument alone explains about 50 percent of the total variation in intermediate relative performance in the last three weeks of the experiment. Figure 5 shows the relation between the actual difference in intermediate cumulative performance between the treatment stores and their best comparison stores, $p^{CU}_{t,w-1} - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$, and the predicted difference (4.6).

Figure 4.5: Actual and predicted difference in intermediate cumulative performance between the treatment store and its best comparison store ($p^{CU}_{t,w-1} - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$)

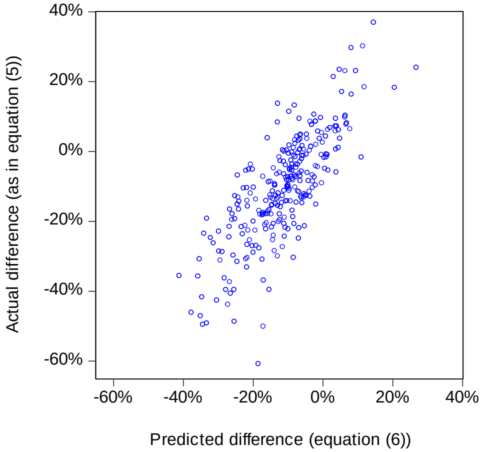

One outlier not shown (actual difference -101%; predicted difference -36%)

The OLS estimation in the first column of Table 3 shows that intermediate relative performance is significantly positively related to subsequent sales. Its point estimate suggests that a percentage point increase in lagged relative performance increases current sales by 0.26 percent. However, in the IV-2SLS estimation reported in the third column the point estimate is more than halved and is no longer significantly different from zero. Figure 6 visualises these results for the relevant subset of

Figure 4.6: The relation between the predicted difference $E(p^{CU}_{t,w-1}) - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$ and the residuals from estimating (3)

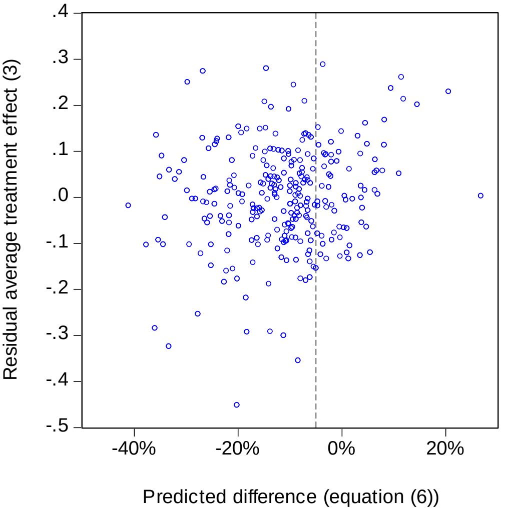

observations: treatment stores in the three final weeks of the experiment. It plots the residuals of the estimation of the average treatment effect (4.3), as presented in the first column of Table 2, against the predicted values for $p^{CU}_{t,w-1} - \max_{c_t}\left[p^{CU}_{c_t,w-1}\right]$ as estimated by the first-stage regression of the IV-2SLS estimation (second column of Table 3).

The fourth column of Table 3 shows that the effect of intermediate relative performance is concentrated in the final week of the experiment. Stores that perform about as good or even better than their best-performing comparison store in the first three weeks of the experiment manage to increase their sales in the final week. A Wald test shows that this increase is significant at the 5 percent level for stores that lag less than 4 percentage points behind their best-performing comparison store. The treatment effect in the final week increases significantly in intermediate relative performance, by 0.28 percent per percentage point.

The estimations in Table 3 assume that the effect of intermediate relative performance is linear. The first column of Table 4 reports the results of estimating (4.8),

Table 4.3: Dynamic Incentives

Table 3: Dynamic incentives

	Dependent variable: ln(sales)							
	All stores				Big stores		Small stores	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	IV-2SLS	IV-2SLS	IV-2SLS	IV-2SLS	IV-2SLS	IV-2SLS	IV-2SLS
Stage IV-2SLS		First	Second	Second	Second	Second	Second	Second
Treatment	0.017	0.023**	0.005		0.028		-0.020	
	(0.012)	(0.010)	(0.014)		(0.019)		(0.020)	
Relative intermediate performance	0.0026***		0.0011		0.0022*		0.0003	
	(0.0005)		(0.0008)		(0.0012)		(0.0011)	
Expected intermediate relative performance, $D_{t,w-1}$ (equation (6))		-1.088***						
		(0.091)						
Treatment week 1				-0.003		0.026		-0.034
				(0.019)		(0.027)		(0.026)
Treatment week 2				0.001		0.033		-0.030
				(0.021)		(0.026)		(0.027)
Treatment week 3				-0.016		0.007		-0.041
				(0.022)		(0.026)		(0.042)
Treatment week 4				0.048**		0.041		0.064**
				(0.022)		(0.031)		(0.032)
Relative performance after week 1				0.0011		0.0049***		-0.0004
				(0.0009)		(0.0016)		(0.0011)
Relative performance after week 2				0.0004		0.0005		0.0002
				(0.0013)		(0.0012)		(0.0024)
Relative performance after week 3				0.0028**		0.0009		0.0051***
				(0.0012)		(0.0018)		(0.0015)
Carnival	-0.020	-0.017**	-0.024	-0.027*	-0.028**	-0.035**	-0.025	-0.027
	(0.014)	(0.008)	(0.015)	(0.016)	(0.014)	(0.015)	(0.021)	(0.023)
Store-fixed effects	yes	yes	yes	yes	yes	yes	yes	yes
Week-fixed effects	yes	yes	yes	yes	yes	yes	yes	yes
Store-week observations	4158	4158	4158	4158	2134	2134	2024	2024
Stores	189	189	189	189	97	97	92	92
R^2	0.9284	0.6712	0.9283	0.9284	0.9050	0.9050	0.8988	0.8990

Standard errors clustered at the store level in parentheses.
The dependent variable in the first-stage regression shown in column (2) is relative intermediate performance.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

Table 4.4: Dynamic incentives separate for stores close to winning the bonus

	Dependent variable: ln(sales)			
	All stores		Big stores	Small stores
	(1)	(2)	(3)	(4)
	IV-2SLS	IV-2SLS	IV-2SLS	IV-2SLS
	Second stage	Second stage	Second stage	Second stage
Treatment	-0.006		0.023	-0.033
	(0.016)		(0.024)	(0.021)
Relative intermediate performance	0.0002		0.0017	-0.0006
	(0.0009)		(0.0017)	(0.0011)
Treatment*close	-0.008		-0.035	0.016
	(0.019)		(0.028)	(0.029)
Relative intermediate performance*close	0.0070***		0.0061***	0.0072***
	(0.0012)		(0.0023)	(0.0012)
Treatment week 1		-0.003		
		(0.019)		
Treatment week 2		-0.045		
		(0.036)		
Treatment week 3		-0.038		
		(0.028)		
Treatment week 4		0.050		
		(0.031)		
Relative performance after week 1		-0.0011		
		(0.0017)		
Relative performance after week 2		-0.0012		
		(0.0016)		
Relative performance after week 3		0.0027		
		(0.0018)		
Treatment week 2*close		0.057		
		(0.042)		
Treatment week 3*close		-0.061		
		(0.055)		
Treatment week 4*close		-0.038		
		(0.038)		
Relative performance after week 1*close		0.0048		
		(0.0030)		
Relative performance after week 2*close		0.0149**		
		(0.0059)		
Relative performance after week 3*close		0.0063**		
		(0.0028)		
Carnival	-0.024	-0.027	-0.026**	-0.027
	(0.015)	(0.017)	(0.013)	(0.021)
Store-fixed effects	yes	yes	yes	yes
Week-fixed effects	yes	yes	yes	yes
Store-week observations	4158	4158	2134	2024
Stores	189	189	97	92
R^2	0.9283	0.9282	0.9051	0.8985

Standard errors clustered at the store level in parentheses.
"Close" is a dummy variable that takes value one when the store's expected intermediate performance is at most 5 percentage points below its best comparison store, i.e. when $D_{t,w-1} > -.05$ (see equation (6)).
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

where the effect of intermediate relative performance is allowed to vary between stores that lag far behind and stores that are close to or above the target for winning a bonus.[13] Graphically, we allow the effect of intermediate relative performance to differ between observations to the left and right of the dashed line in Figure 6. We find that past relative performance does not affect current sales for stores that lag far behind. In contrast, current sales of treatment stores that lag less than 5 percentage points behind increases by 0.72 percent per percentage point increase in past relative performance. A Wald test shows that the overall treatment effect is significantly different from zero for stores that are at least 6 percentage points ahead of their best comparison store. The second column of Table 4 shows that the marginal effect of intermediate relative performance on current sales of relatively good-performing stores is significantly positive in the third and fourth week of the experiment, with magnitudes of 1.4 and 0.9 percent per percentage point, respectively. For stores that lag far behind, there is no such effect in any of the weeks. These results are qualitatively robust to varying the level of intermediate relative performance at which stores are deemed to stand a chance of winning between 0% and -10%. Quantitatively, the estimated effects of intermediate relative performance on current sales for stores deemed to stand a chance are larger when this level is closer to 0%.

Lastly, we examine differences in treatment effects between big and small stores. This is explorative, as store size is not exogenously determined. Hence, any differences between big and small stores may not be caused by store size per se, but by an unobserved store characteristic that causes or co-varies with store size. Unreported estimations show that we do not find a significantly positive average treatment effect in any of the strata based on store size. Columns 5 to 8 of Table 3 show the results of estimating the effect of intermediate relative performance (4.5) for big and small stores separately, both for the whole period as well as separated by week. For big stores, the treatment effect increases marginally significantly with intermediate rela-

[13] Instead of estimating (4.8), we could estimate a quadratic specification of intermediate relative performance. However, the estimates for the quadratic specification would be heavily affected by the many treatment store-week observations with sizable negative intermediate relative performance (see Figure 5). Hence, we would learn little about the marginal effect of intermediate relative performance for stores close to winning a bonus.

tive performance. Column 6 shows that most of this effect occurs in the week where stores receive the first poster (i.e. in the second week of the experiment). For small stores, we find a sizeable treatment effect in the last week of the experiment, which increases significantly in intermediate relative performance. The point estimate of the treatment effect in the final week for small stores that perform at par with their best comparison store up to the final week is 6.4% additional sales. Columns 3 and 4 of Table 4 show that both big and small stores respond strongly to intermediate relative performance when sufficiently close to their best comparison store. For big stores, the estimated effect of a percentage point increase in intermediate relative performance on current sales is 0.78 percent. For small stores, this is 0.66 percent. Wald tests show that both effects are statistically significant with p-values < 0.01.[14]

Taken together, these results paint the following picture. On average, the relative performance incentive scheme had no effect on sales. Possibly, the prospect of competing against non-competitors did not excite employees in the treatment stores much. Alternatively, many stores may have perceived the relative performance targets as too ambitious. Such a perception would be reinforced after receiving the first poster with rankings, as only 23 treatment stores ranked on top of the first poster and 64 stores lagged more than 5 percent behind their best-performing comparison store on the first poster. We find that stores lagging too far behind do not respond to the incentive scheme, nor to the intermediate relative performance information. However, as stores come closer to winning a bonus through better lagged relative performance, sales increase significantly with lagged relative performance. This effect is strongest in the second half of the experiment, and is present in both big and small stores. This result contrasts with Casas-Arce and Martinez-Jerez (2009), who do not find that performance decreases when trailing contestants lag further behind. One explanation, as conjectured by the authors, is that attrition bias may affect their result, while it is absent in our study. Frank and Obloj (2011) also study a competition without attrition and find, like us, that performance is increasing in

[14] Across the three weeks with intermediate relative performance figures, there are 19 respectively 33 observations of small respectively big stores whose performance is at most 5 percent below their best comparison store. This implies that we cannot further disaggregate the results in columns 3 and 4 of Table 4 into estimates by week.

intermediate relative performance for contestants that lag behind.

4.5 Concluding remarks

We have reported the results of a field experiment on dynamic incentive effects of relative performance pay among stores of a retail chain. We find that intermediate relative performance feedback affects subsequent performance of stores close to the bonus target. These stores show significantly higher performance, particularly near the end of the incentive period. Stores lagging far behind do not respond to the incentive scheme, nor to intermediate relative performance. As many treatment stores happen to trail far behind bonus targets over the course of the experiment, we find no improvement in performance on average.

Chapter 5

The Effects of Prize Spread and Noise in Elimination Tournaments: A Natural Field Experiment

Joint with Josse Delfgaauw, Robert Dur, and Willem Verbeke.

5.1 Introduction

Tournament theory is a cornerstone of incentive theory in organizations. Pioneered by Lazear and Rosen (1981), Green and Stokey (1983), Nalebuff and Stiglitz (1983), and Rosen (1986), tournament theory can explain many prominent organizational features. Examples include large wage increases upon promotions (as found by e.g. Murphy 1985, Baker et al. 1994, McCue 1996), a convex wage structure across the levels of the hierarchy (Murphy 1985, Baker et al. 1994, Gibbs 1995), and a positive relation between the prize spread and the number of people competing for a promotion (Eriksson 1999, Bognanno 2001). Waldman (2008) provides a more extensive discussion of empirical evidence on tournament theory. Crucially, predictions from tournament theory for organizational architecture follow from employees' presumed responses to tournament incentives.

In this chapter, I report the results of a natural field experiment we conducted in a privately-held company. We design elimination tournaments with two rounds that allow us to test several basic hypotheses on employees' behavior as derived from standard tournament theory. First, we vary the distribution of total prize money over the two rounds of the elimination tournament. Theory predicts that a more convex prize structure while keeping total prize money constant (i.e. simultaneously decreasing the prize for winning the first round and increasing the prize for winning the second round) leads to better second-round performance at the expense of first-round performance. Second, we investigate whether the level of noise in contestants' performance affects their performance in the tournament. In theory, noise dilutes incentives to perform, as it reduces the marginal effect of effort on the probability of winning.[1]

To test these hypotheses, we run an elimination tournament among a randomly chosen subset of the 208 stores of a retail chain in the Netherlands selling music, movies, and video games. Both rounds of the tournament last four weeks. Performance in the tournament is measured by the Average number of Products per Customer (APC), a relatively stable and well-known performance measure in this company. In the first round, the 144 participating stores are assigned to poules of four stores that are comparable in terms of historical performance. After the first round, the two worst-performing stores of each poule are eliminated, whereas the two best-performing stores of each poule win a prize and proceed to the second round. In the second round, stores are once more grouped into poules of four comparable stores. The two best-performing stores of each poule again win a prize.

To investigate the relation between prize structure and the incentive effect of the tournament, participating stores are assigned to two different treatments. The treatments differ by the prize spread only, we keep the total amount of prize money constant. In the low-spread treatment, prizes are identical in the two rounds, whereas in the high-spread treatment the second-round prize is four times as large as the first-round prize.

[1] Our design allows for a clean test of the effects of prize spread and noise on employees' incentives to perform well. Tournament theory also generates predictions on the effects of participant heterogeneity (Lazear and Rosen 1981, Rosen 1986), the incentives to sabotage (Lazear 1989), the choice of low-risk versus high-risk strategies (Knoeber and Thurman 1994), and self-selection into tournaments (Lazear and Rosen 1981). See Charness and Kuhn (2010) and Lazear and Oyer (2009) for recent overviews.

Our findings are largely in line with theoretical predictions. First, the average incentive effect of the tournament is approximately 1.5 percent. This effect is statistically significant. Economically, the effect is not particularly large, but still relevant given the nature of the retail chain's business. Second, we find that second-round performance is 1 percentage point higher in the high-spread treatment as compared to the low-spread treatment, while first-round performance is 0.8 percentage point lower. These differences are qualitatively in line with theory, but they are not statistically significant. Third, in the high-spread treatment, the estimated second-round treatment effect is significantly higher than the first-round effect, as predicted by theory. Fourth, while theory predicts a higher first-round performance as compared to second-round performance in the low-spread treatment, we find the reverse, albeit insignificantly so. As a result, most of the average treatment effect is concentrated in the second round of the tournament.

To test for the effect of noise in measured performance on the effect of the tournament, we use the variance in performance prior to the experimental period as our measure of noise. In the assignment of stores to poules, we take their level of noise into account, so that store with relatively low (high) noise in performance are matched to other stores with relatively low (high) noise. As predicted by theory, noise has a negative effect on the response to the tournament. This effect is mainly concentrated in the second-period. The impact is substantial relative to the average treatment effect: while the stores with least noise experience an estimated treatment effect of about 2.4 percent, the estimated treatment effect is zero for the quartile of stores with highest noise in performance.

The remainder of this chapter is organized as follows. The next section discusses related empirical work. The design of the experiment is discussed in Section 5.3. In Section 5.4, we analyze a simple elimination tournament model to derive five hypotheses that our experiment allows us to test. Section 5.5 provides summary statistics and Section 5.6 describes our estimation strategy. Our findings are presented and discussed in Section 5.7.

5.2 Previous studies

Two studies use non-experimental field data to test similar hypotheses from tournament theory. Audas et al. (2004) use the administrative records of a British financial firm to investigate the effects of prize spread and noise in promotion decisions on the absence rate of employees. They find that larger prize spreads, defined as the difference in average earnings between two adjacent layers in the firm's hierarchy, reduce absenteeism. More unexplained variation in promotion decisions increases absenteeism. Based on data from a cross-section of firms, DeVaro (2006) estimates a structural model treating prize spread, performance, and promotions as endogenous. He finds a positive effect of prize spread on workers' performance ratings, a negative effect of noise on performance, and a positive effect of noise on the prize spread. We see our methodology as complementary. By conducting a field experiment rather than analysing actual career paths, we can generate exogenous variation in prize spread and obtain a simple measure of noise in performance. This allows for an easy identification of the effects of prize spread and noise on performance in tournaments within an organization.[2]

Field experiments in organizations are scarce. To our knowledge, this is the first field experiment that studies the effects of an elimination tournament. Field experiments with one-stage tournaments have been conducted by Erev et al. (1993) and Bandiera et al. (2009) among fruit pickers and by Casas-Arce and Martinez-Jerez (2009) and Delfgaauw et al. (2009) among retailers. These studies find a positive effect of tournament incentives on performance, but none of them varies the prize spread. Lim et al. (2009) vary both the number and the distribution of prizes in contests among fundraisers, keeping total prize money constant. They find that contests with multiple identical prizes elicit higher effort as compared to single-prize contests, but differentiating prizes by rank has no effect on effort. None of these studies looks at the effect of noise in the performance measure.

[2] Several studies test elements of tournament theory in a non-organizational setting. Ehrenberg and Bognanno (1990) show that golf players' performance increases in the effect of improvements in relative positions on prize money. Becker and Huselid (1992) find that race car drivers' performance increases in prize spread. Knoeber and Thurman (1994) find a similar result in competitions among broiler producers.

In terms of design, our study is closely related to a recent laboratory experiment by Altmann et al. (2008). In a chosen-effort setting, they find that subjects choose significantly higher effort in the first stage of a two-stage elimination tournament as compared to a strategically equivalent one-stage tournament. A more convex prize spread in the elimination tournament, obtained by decreasing the prize for winning the first round, does not affect effort in either stage of the tournament, in contrast to theory. Sheremeta (2010) and Höchtl et al. (2011) compare a one-shot contest with a multi-stage contest for the case of a single prize, using contest success functions to determine the winner. Whereas parameters are set such that total effort should be equivalent in the two treatments, both studies find substantial overprovision of effort in the elimination tournament relative to the one-shot contest. Sheremeta (2010) shows that the amount of overprovision relates to the level of effort subjects chose in a contest without a monetary prize, suggesting that some people experience non-monetary utility from winning.[3]

The effects of noise on performance in tournaments are rarely studied in experiments. An exception is Bull et al. (1987), who find in a laboratory experiment with chosen-effort that a simultaneous change in noise and marginal cost of effort that leaves equilibrium effort unaffected indeed leads to similar levels of effort chosen by subjects. Given the prevalence of relative performance incentives and noisy performance measures in real-world settings, our study provides an important test of this part of tournament theory.

5.3 Experimental design

The experiment took place in the period September - November 2010 in a retail chain in the Netherlands that sells computer games, music, and movies. In September 2010, the retail chain owned 208 geographically dispersed stores, operating under two different brands.

We set up an elimination tournament for a randomly selected set of stores. Per-

[3] Several lab experiments analyse the effects of prize spread in a standard tournament setting, see e.g. Bull et al. (1987), Harbring and Ihrlenbusch (2003), and Freeman and Gelber (2010). Charness and Kuhn (2010) provide an overview.

Figure 5.1: Mean of Average number of Products per Customer (APC) across all stores, by week

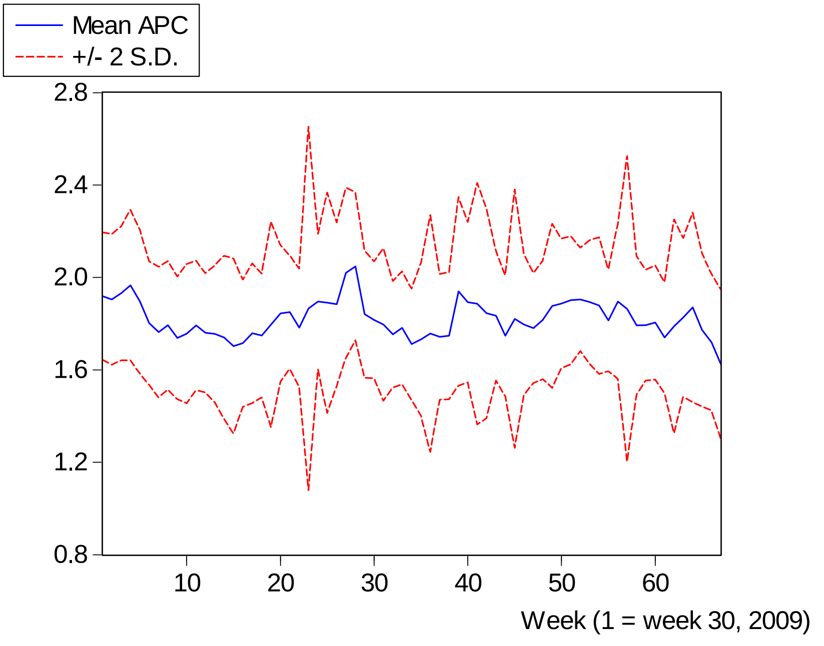

formance in the tournament was based on the Average number of Products sold per Customer (APC). This performance measure can be directly computed from the company's operational database, which records the number of products sold per store per week, and the number of customers (i.e. transactions) per store per week. Everyone in the company is familiar with APC as performance measure. It is part of employees' standard incentive scheme, and stores receive a weekly report on their performance including APC. The reason that APC was chosen as a performance measure and not, for instance, sales is twofold. First, it makes unequivocally clear how stores can enhance performance: through an increase in cross-selling. Second, there is relatively little variation in this measure over time. Figure 1 shows the average APC per week for the period of week 30 in 2009 until week 45 in 2010.

The elimination tournament consisted of 2 rounds, both lasting four weeks with a two-week break in between. The first round ran from week 36 to week 39 in 2010, the second round from week 42 to week 45. In the first round, the participating stores were assigned to poules of four stores. All employees of the two best-performing stores per poule, i.e. those with highest APC cumulative over the four weeks in round 1, received a bonus. Moreover, these stores qualified for the second round of the tournament, while all other stores (the bottom-two stores of each poule) were eliminated. In the second round, qualified stores were again assigned to poules of four. As before, all employees of the top-two performing stores per poule over the second round received a bonus. After round 2, the tournament ended.

We scheduled a two-week gap between the end of the first round and the start of the second round. This period was used to communicate the results of the first round to all treatment stores and to inform the winners of the first round of their second-round assignment. This two-week period is not included in the estimations below, as otherwise a response to winning or losing would affect the estimates of the store-fixed effects.

In February 2010, we ran another experiment at the same retail chain, aimed at studying dynamic incentive effects of tournaments (the results are reported in Delfgaauw et al., 2010). At that time, a randomly selected set of stores could earn an additional bonus, while the remaining stores were promised a similar opportunity later in the year. Hence, all stores that did not participate in the first experiment (115 stores) do participate in the current elimination tournament. Furthermore, to check whether assignment to treatment or control in the first experiment affects performance in the elimination tournament, we randomly select an additional 29 stores from the stores that comprised the treatment group in the first experiment to participate also in the current tournament. Below, we check whether these 29 stores respond differently to the current treatment as compared to the stores that were part of the control group in the February 2010 experiment. If we find no differences, we can be assured that there are no spill-over effects of the February 2010 experiment. In that experiment, a total of 15 stores were (non-randomly) not allowed to participate in the first experiment, for a variety of reasons. One of these stores was closed during 2010. Furthermore, 6 new stores were opened during the year. These 20 stores all participate in the current tournament, but since they were non-randomly assigned, they are left out of the analysis. Furthermore, two stores were not allowed to participate in the current experiment and, hence, are also left out of the analysis. This leaves us with 186 stores in the analysis. Of these stores, 62 comprise the control group, while the remaining 124 comprise the treatment group.

To study the effect of prize spread on the incentive effect of the tournament, we assign the participating stores to two different treatment groups. The only difference between the two treatments is the prize spread. Thus, we keep total prize money identical across the treatments. In the first round of the low-spread treatment, the bonus for being one of the two best-performing stores in the poule is 35 euro gross for a full-time employee. In the second round, the bonus is again 35 euro gross. Hence, per eight stores, employees of two stores win in total 70 euro, employees of two other stores win 35 euro, while the employees of four stores win nothing. In the high-spread treatment, the first-round bonus is 17.50 euro gross. The bonus in the second round is 70 euro. Hence, per eight stores, employees of two stores in the high-spread treatment earn 87.50, employees of two other stores earn 17.50, and four stores receive nothing. Comparing the two treatments, the expected monetary bonus per employee at the start of the tournament is identical in both treatments (26.25 euro), while the prize spread is higher in the high-spread treatment than in the low-spread treatment.[4] All bonuses earned were paid after the tournament ended (in December).

We also examine the effect of noise in performance on the incentive effect of the tournament. We take stores' standard deviation in the performance measure APC over the period August 2009 to August 2010 as our measure of noise. Note that this period does not include the experimental period, so that our measure of noise is not affected by the response to the tournament incentives. Furthermore, store's assignment to poules was partially based on this performance measure, as described below, so that high-noise (low-noise) stores competed against other high-noise (low-noise) stores.

The assignment of stores in the treatment group to the different treatment condi-

[4] Parttimers receive an amount proportional to their contract size.

tions (low prize spread and high prize spread) went as follows. First, we stratified the stores by their noise in the performance measure. We divided them in two equally large groups, one group with the stores with the highest standard deviation in APC and one group with stores with the lowest standard deviation. Subsequently, we randomly assigned half of the stores in each noise-group to the low-spread treatment and the other half to the high-spread treatment. By doing so, we created four groups of equal size (31 stores) that differ in two dimensions: high noise stores versus low noise stores, and low-spread treatment versus high-spread treatment. A similar procedure was used to assign the 20 non-randomly selected stores to these four groups, so that each group contains 36 stores. In the first round of the tournament, stores compete against three other stores from the same group. The assignment to poules is based on historical performance. Per group, we rank stores on average performance (APC) in the period August 2009 to August 2010. The best-scoring four stores are placed together in a poule, as well as numbers 5 to 8, etc. This creates in total 36 poules, with 9 poules for each treatment-noise group.

In the second round, we again assigned stores to poules on the basis of average performance (APC) in the period August 2009 to August 2010. In both treatments, we kept the stores in the high-noise group and the low-noise group separate, with one exception: in both treatments, the two stores with the lowest APC in each of the two noise-groups were placed together in a poule.[5] Hence, in the second stage of the tournament, we have in total 72 stores divided over 18 poules, with 4 poules per treatment-noise group plus 1 poule per treatment with stores from both the low-noise group and the high-noise group. Out of the 20 non-randomly assigned stores participating in the tournament, 9 made it to the second round. Hence, we can use 63 participating stores in the analysis of the second-round treatment effects, which are almost equally divided over the four treatment-noise groups.

All communication about the elimination tournament to the stores went through the company's internal communication channel. Stores were not aware of our in-

[5] As it turns out, seven of these 8 stores were among the 20 stores that were non-randomly assigned to a treatment and are therefore left out of the analysis. The remaining store belongs to the high-spread, high-noise group. We treat this store the same as all other stores in this group. Leaving the store out of the analysis does not affect the results.

volvement, so that our experiment qualifies as a natural field experiment (Harrison and List 2006). A week before the first round started, all stores of the retail chain learned that a new incentive event would take place. We explained that all stores who did not participate in the February 2010 experiment would participate in the current experiment, as well as a randomly selected number of stores that did participate in February.[6] A few days later, all participating stores received a message with the rules of the elimination tournament. Stores in the high-spread and low-spread treatment received identical messages, except for the amounts of prize money mentioned for winning the first and second round. We informed them that other stores, randomly selected, faced a different division of prize money, to reduce confusion and suspicion that might arise when employees learn during the tournament that other stores were entitled to different prizes. Also, we explained that assignment to poules in the second round would be based on the average APC over the period of August 2009 up to August 2010, not on performance during the first round, so as to avoid ratchet effects. Just before the start of each round, the stores (still) participating in the tournament received the assignment to poules for all stores, with for each store the average APC over the period of August 2009 up to August 2010. Hence, the stores could verify that they were matched to stores with similar historical performance. During the tournament, each store received weekly feedback on the ranking of stores in its poule in the form of a poster with APC-figures for all stores in the poule. These posters could be attached to a larger poster, which store managers were instructed to hang in a prominent place (typically the store's canteen).

5.4 Deriving hypotheses

In this section, we set up a simple model to derive the hypotheses that our experiment allows us to test. For a general treatment of the effects of prize spread and noise in tournaments, see Lazear and Rosen (1981) and Gibbs (1996). A general model of incentive effects of elimination tournaments can be found in Rosen (1986).[7]

[6] As few stores actually won a bonus in the February-experiment, this was not perceived as unfair.

[7] Recent theoretical advances on elimination tournaments include endogenizing the number of stages and the prize structure (Fu and Lu 2009) and optimal seeding when participants are

Consider four identical agents that participate in a two-stage elimination tournament. In the first stage, the agents compete pairwise. The winners of the first stage receive prize $B_1 \geq 0$ and go on to the second stage of the tournament. The first-stage losers are eliminated from the tournament and receive nothing. In the second stage, the two first-stage winners compete against each other for one prize with value $B_2 > 0$.[8]

Let $Q_{i,t}$ be agent i's performance in stage t. Performance depends on effort $e_{i,t}$ and on idiosyncratic noise $\mu_{i,t}$:

$$Q_{i,t} = q(e_{i,t}) + \mu_{i,t},$$

where $q(\cdot)$ is concave. Effort and noise are not observable, performance is verifiable. Agent i's utility in stage t depends on income $w_{i,t}$ and effort cost:

$$U_{i,t} = w_{i,t} - c(e_{i,t}),$$

where $c(\cdot)$ is strictly convex: $c' > 0$, $c'' > 0$. We neglect discounting across stages of the tournament and assume an interior solution for optimal effort. We aim to derive a symmetric subgame-perfect Nash equilibrium.

In the contest between agents i and j in stage t, let $\Delta\mu_{ij,t} = \mu_{i,t} - \mu_{j,t}$ be the noise difference. We assume that $\Delta\mu_{ij,t}$ is distributed according to density function $f(\cdot)$ which is unimodal and symmetric around zero and twice continuously differentiable, with cumulative density function $F(\cdot)$. Across contests, draws of $\Delta\mu_{ij,t}$ are independent. Given effort $\hat{e}$ of contender j in a given stage, agent i's probability of winning that stage is given by $p[q(e_{i,t}) - q(\hat{e}) > -\Delta\mu_{ij,t}] = 1 - F[q(\hat{e}) - q(e_{i,t})]$. Hence, the marginal effect of effort on the winning probability is given by $f[q(\hat{e}) - q(e_{i,t})]q'(e_{i,t})$.

First, consider the second stage. In a symmetric equilibrium, both agents opti-

heterogeneous (Groh et al. 2010).

[8] In the experiment, we have competition between teams rather than between individuals. Also, rather than competition in poules of 2, we have 4 contestants per poule competing for two prizes per poule. This does not qualitatively affect the predictions from the theoretical analysis.

mally exert the same level of effort, as implicitly given by first-order condition

$$f(0)q'(e^*_{t=2})B_2 - c'(e^*_{t=2}) = 0. \tag{5.1}$$

In the symmetric equilibrium, the probability of winning the second stage is equal to $F(0) = \frac{1}{2}$, so that second-stage expected utility (conditional on winning the first stage) equals $U_{t=2} = \frac{1}{2}B_2 - c(e^*_{t=2})$. As a result, the expected value of winning the contest in the first stage is given by $B_1 + \frac{1}{2}B_2 - c(e^*_{t=2})$. In a symmetric equilibrium, maximising first-stage utility yields the following first-order condition for optimal effort

$$f(0)q'(e^*_{t=1})\left[B_1 + \frac{1}{2}B_2 - c(e^*_{t=2})\right] - c'(e^*_{t=1}) = 0. \tag{5.2}$$

By applying the implicit function theorem to first-order conditions (5.2) and (5.1), we derive the following predictions regarding the effects of noise in the performance measure and of the prize structure on performance in the elimination tournament. Proposition 1 describes the effect of noise.[9]

Proposition 5.1 *A larger variance of the noise distribution $f(\cdot)$, so that mass is shifted from the mode to the tails, leads to lower performance in both stages of the tournament.*

Proof. *Higher variance reduces $f(0)$. Totally differentiating (5.1) gives*

$$\frac{\partial e^*_{t=2}}{\partial f(0)} = -\frac{q'(e^*_{t=2})B_2}{f(0)q''(e^*_{t=2}) - c''(e^*_{t=2})} > 0.$$

The effect on first-round effort is similar. ■

Next, we derive the effects of increasing the convexity of the prize spread. Consider two tournaments with identical total prize money, but different prize spreads. Using superscript L (H) to refer to the tournament with low (high) prize spread, we

[9] In estimating the effects of noise, we use the variance in individual stores' performance μ_i rather than the variance of the difference in the error terms $\Delta\mu_{i,j}$ as in Proposition 1. This has qualitatively no effect on the hypothesis of the effect of noise. The distribution of the difference between two i.i.d. random variables with density $f(\cdot)$ is unimodal with a maximum at zero when $f(\cdot)$ is unimodal (Vogt 1983). By Bienaymé's formula, the variance of the difference of two i.i.d. random variables is the sum of the variance of the two variables. Hence, the variance of $\Delta\mu_{i,j}$ is increasing in the variance of μ_i.

have $B_1^L > B_1^H$, $B_2^L < B_2^H$, and $2B_1^L + B_2^L = 2B_1^H + B_2^H$. This yields the following predictions regarding the effect of prize spread on performance in the second and first round, respectively.

Proposition 5.2 *Second-stage performance in the high-spread tournament is better than second-stage performance in the low-spread tournament.*

Proof. *Totally differentiating (5.1) shows that $e^*_{t=2}$ increases in B_2:*

$$\frac{\partial e^*_{t=2}}{\partial B_2} = -\frac{f(0)q'(e^*_{t=2})}{f(0)q''(e^*_{t=2}) - c''(e^*_{t=2})} > 0.$$

$B_2^H > B_2^L$ *implies that* $e^H_{t=2} > e^L_{t=2}$. ■

Proposition 5.3 *First-stage performance in the low-spread tournament is better than first-stage performance in the high-spread tournament.*

Proof. *By Proposition 2 and $B_2^L < B_2^H$, second-stage effort is higher in the high-spread treatment, so that $c(e^L_{t=2}) < c(e^H_{t=2})$. As total prize money is identical, we have $B_1^L + \frac{1}{2}B_2^L = B_1^H + \frac{1}{2}B_2^H$, so that $B_1^L + \frac{1}{2}B_2^L - c(e^L_{t=2}) > B_1^H + \frac{1}{2}B_2^H - c(e^H_{t=2})$. By (5.2), this implies that $e^L_{t=1} > e^H_{t=1}$.* ■

Propositions 2 and 3 together imply that a higher prize spread increases second-round performance at the expense of first-round performance. A higher second-stage bonus induces agents to exert more effort in the second round, which reduces the expected value of winning the first round.

Lastly, given a certain prize structure, the model provides predictions on first-round performance in the tournament relative to second-round performance.

Proposition 5.4 *If $B_1 \geq B_2$, performance in the first stage is better than performance in the second stage.*

Proof. *Second-stage utility $\frac{1}{2}B_2 - c(e^*_{t=2}) > 0$. Hence, if $B_1 \geq B_2$, $B_1 + \frac{1}{2}B_2 - c(e^*_{t=2}) > B_2$. Comparing (5.1) and (5.2), it follows that $e^*_{t=1} > e^*_{t=2}$.* ■

Proposition 5.5 *If $B_1 \leq \frac{1}{2}B_2$, performance in the first stage is worse than performance in the second stage.*

Proof. *Second-stage effort cost $c(e^*_{t=2}) > 0$. Hence, if $B_1 \leq \frac{1}{2}B_2$, $B_1 + \frac{1}{2}B_2 - c(e^*_{t=2}) < B_2$. Comparing (5.1) and (5.2), it follows that $e^*_{t=1} < e^*_{t=2}$.* ■

In the experiment, we have two treatments with identical total prize money but different prize structures. The first treatment has a relatively low prize spread, with equal prizes in both rounds: $B_1^L = B_2^L$. The second treatment has a relatively high prize spread, with $B_1^H = \frac{1}{4}B_2^H$. Hence, Propositions 2 and 3 predict that stores in the low-spread treatment show better first-round performance than stores in the high-spread treatment, but lower second-round performance. Furthermore, Proposition 4 predicts that in the low-spread treatment the first-round treatment effect should be higher than the second-round treatment effect, while Proposition 5 predicts the reverse for the high-spread treatment. Lastly, in both treatments we divide the stores in two groups depending on the historical variance of the performance measure. Proposition 1 predicts that we should observe lower treatment effects among the stores in the high-noise group compared to the stores in the low-noise group.

5.5 Summary statistics

In our estimations, performance is given by the Average number of Products per Customer per week (APC). Table 1 shows that on average, a customer buys 1.82 products per transaction. During the experimental period, the average APC-score is somewhat lower than in the year preceding the tournament. Comparing the stores in the control group with the stores in the high and low prize spread group, we find no differences in historical performance. APC is a relatively stable performance measure. Averaged across stores, the within-store standard deviation over the period August 2009 to August 2010 is 0.15. There is some variation in this measure of noise across stores, as it ranges from a minimum of 0.07 to a maximum of 0.54, with a median of 0.13. Figure 2 shows that the distribution of noise is very similar across the control group and the high-spread and low-spread treatment groups.[10] In other observable store characteristics, we find no statistically significant differences except for the share of female employees: stores in the control group have relatively few female employees.

Grouping the treatment stores by noise group, we find that treatment stores

[10]Figure 2 shows that there are a few stores with unusually large standard deviations in APC. None of the results in this chapter are affected by removing these stores from the analysis.

with a large standard deviation in APC show a higher average APC, which is an indication of heteroscedasticity. The difference in noisiness of the performance measure between the low-noise and the high-noise group is substantial. The standard deviation of APC in the high-noise group is about 50% larger than in the low-noise group. Proposition 1 states that the treatment effect should be decreasing in the variance of the performance measure, provided that the density at the mode of the error distribution is smaller for high-noise stores than for low-noise stores. Figure 3 suggests that this holds in the data, by showing kernel densities of the residuals of a regression of APC on store-fixed effects using all observations before the tournament starts. The peak of the kernel density is clearly lower for stores in the high-noise group than for stores in the low-noise group. In both groups, the peak lies marginally below zero. Other store characteristics show no differences between the high and low-noise stores.

5.6 Estimation

We assess the effects of the tournament incentives using OLS with week fixed effects and store fixed effects. Let $y_{i,w}$ be the performance of store i in week w. Let T_i (C_i) be a dummy variable that takes values 1 for treatment (control) stores. Furthermore, based on the results of the first round of the tournament we create a dummy W_i that takes value 1 for stores that have won in the first round (and, hence, take part in the second round of the tournament) and a dummy E_i that takes value 1 for the stores that are eliminated from the tournament after the first round. Lastly, R_1 and R_2 are two dummy variables indicating the weeks in which the first and second round of the tournament took place, respectively. We estimate the average treatment effect by

$$y_{i,w} = \alpha_i + \tau_w + \beta T_i \left[R_1 + W_i R_2 \right] + \delta E_i R_2 + \varepsilon_{i,w} \tag{5.3}$$

where α_i and τ_w are store and week fixed effects, respectively, and $\varepsilon_{i,w}$ is an error term.[11] Coefficient β gives the average treatment effect of stores in competition

[11] In our estimations we cluster standard errors at the store level to correct for serial correlation within stores and heteroscedasticity across stores, as suggested by Bertrand et al. (2004).

Table 5.1: Descriptive Statistics

Table 1: Descriptive statistics

	All stores		Control group		Treatment group: Low-spread		Treatment group: High-spread		Noise group: Low noise		Noise group: High noise	
	mean	Std	mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
Performance (store averages)												
APC, all weeks**	1.82	0.10	1.81	0.09	1.82	0.12	1.83	0.08	1.81	0.08	1.85	0.12
APC assignment period (weeks 32/2009 - 30/2010)**	1.83	0.10	1.82	0.09	1.83	0.12	1.83	0.08	1.81	0.08	1.85	0.12
APC Round 1 (weeks 36/2010 - 39/2010)	1.78	0.10	1.77	0.09	1.80	0.12	1.79	0.08	1.79	0.10	1.80	0.10
APC Round 2 (weeks 42/2010 - 45/2010)	1.75	0.13	1.73	0.11	1.76	0.15	1.76	0.12	1.76	0.13	1.75	0.14
APC Round 2, first-round winners					1.79	0.18	1.80	0.11	1.79	0.15	1.80	0.15
APC Round 2, first-round losers					1.72	0.11	1.72	0.11	1.73	0.11	1.71	0.11
Noise												
Within-store standard deviation of APC (noise)+++ in the assignment period (weeks 32/2009 - 30/2010)	0.15	0.06	0.15	0.07	0.14	0.04	0.15	0.07	0.12	0.01	0.18	0.06
Store characteristics												
Number of employees	5.03	2.17	4.93	1.82	5.00	2.11	5.15	2.55	5.11	2.55	5.03	2.11
Percentage female employees**	0.36	0.25	0.31	0.25	0.42	0.25	0.36	0.25	0.39	0.25	0.39	0.25
Mean tenure of employees (months)	38.09	27.96	35.11	31.06	38.55	29.90	40.56	22.37	38.68	28.83	40.42	23.74
Mean age of employees	24.49	3.68	24.41	3.91	24.32	3.76	24.74	3.40	24.86	4.05	24.20	3.02
Number of stores	186		62		62		62		62		62	

APC is defined as Average number of Products per Customer

For one store in the control group, store characteristics were not available.

***, **, * denote statistically significant differences at the 1%, 5%, and 10% level, respectively, between control, low-spread and high-spread stores.

+++, ++, + denote statistically significant differences at the 1%, 5%, and 10% level, respectively, between control, low-noise and high-noise stores.

Figure 5.2: Kernel densities of within-store standard deviation of APC in the assignment period (August 2009 - August 2010), by treatment group

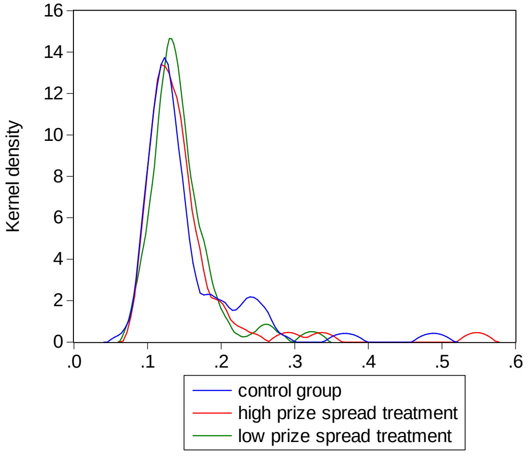

Figure 5.3: Kernel densities of the residuals of regressing APC on store-fixed effects, by noise-group

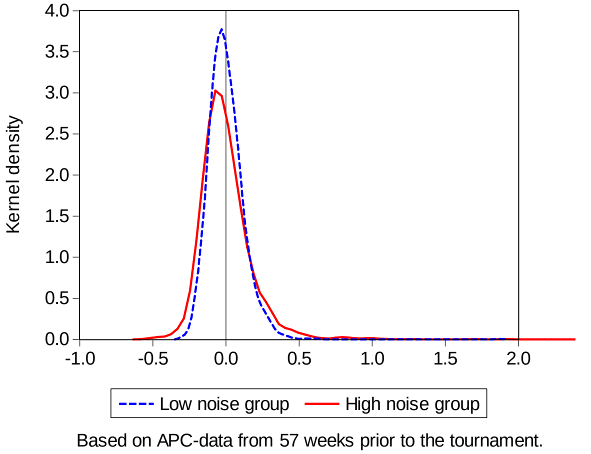

Based on APC-data from 57 weeks prior to the tournament. The right tail of the kernel density of the high-noise group is cut of at 2 for reasons of visibility. It runs to 3.6, based on 6 observations between 2 and 3.6.

versus the stores in the control group. The stores that lost in the first round are non-randomly selected and may respond to losing. Hence, these stores cannot be used as control stores in the second round, and therefore we separate out their second-round performance from the estimation of the treatment effect. It is straightforward to adjust (5.3) to separate the first and second round average treatment effect.

To estimate how the level of noise in a store's performance measure affects the response to the tournament incentives, we use the standard deviation in the performance measure APC over the period August 2009 to August 2010 as a measure of noise. By interacting the treatment dummy with the standard deviation sd_i, we can assess whether the treatment effect is heterogenous in noise, as predicted by Proposition 1. This implies estimating

$$\begin{aligned} y_{i,w} &= \alpha_i + \tau_w + \beta T_i \left[R_1 + W_i R_2 \right] + \nu T_i \left[R_1 + W_i R_2 \right] sd_i + \\ &\quad \mu \left[R_1 + \left(W_i + C_i \right) R_2 \right] sd_i + \delta E_i R_2 + \varepsilon_{i,w} \end{aligned} \tag{5.4}$$

where ν measures how sensitive the treatment effect is to noise, and the term $\mu \left[R_1 + \left(W_i + C_i \right) R_2 \right] sd_i$ measures how performance during the experimental period relates to the standard deviation in APC. The latter term is necessary to control for any time-specific effects of noise, which might otherwise be picked up by ν. Note that we continue to leave out the first-round losers from the estimation of the second-round effects.

To estimate the effect of prize spread, we split dummy T_i into two treatment group dummies. Variable T_L (T_H) takes value 1 when store i is assigned to the low-spread (high-spread) treatment. Replacing T_i in (5.3) by the two treatment group dummies gives

$$y_{i,w} = \alpha_i + \tau_w + \left[\beta_L T_L + \beta_H T_H \right] \left[R_1 + W_i R_2 \right] + \delta E_i R_2 + \varepsilon_{i,w} \tag{5.5}$$

Again, this expression is easily adjusted to estimate the treatment effects in the two tournament rounds separately.

Table 5.2: Average Treatment Effect

	Dependent variable: APC	
	(1)	(2)
Treatment * (Round 1 + Winners * Round 2)	0.028	
	(0.011)**	
Treatment * Round 1		0.014
		(0.011)
Winners * Round 2		0.047
		(0.018)***
Losers * Round 2	-0.014	-0.005
	(0.015)	(0.015)
Store-fixed effects	yes	yes
Week-fixed effects	yes	yes
Store-week observations	12079	12079
Stores	186	186
R^2	0.4471	0.4473

Standard errors clustered at the store level in parentheses.
Winners and Losers refer to the outcome of the first round of the tournament for the treatment stores.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

5.7 Results

5.7.1 Average treatment effect

The first column of Table 2 gives the results of estimating (5.3). We find a statistically significant effect of the tournament on performance. The average treatment effect is 0.028 extra products per customer. This corresponds to an increase of 1.5% of the mean score on Average number of Products sold per Client (APC) and to 20% of within-store standard deviation of APC. The second column of Table 2 separates the estimated average treatment effect by tournament round. On average, the first-round effect is positive but statistically insignificant. In the second round, the treatment effect is about 2.5 percent extra products per customer. The difference between the estimates for the first and second-round treatment effect is statistically significant with a p-value of 0.056. Both estimations show that the stores that lost in the first round perform about as well as the stores in the control group during the second-round period. Hence, two weeks after their elimination, these stores seem to have returned to business-as-usual performance.

We want to exclude that the higher second-round treatment effect is caused by differences in time trends across stores. If some stores experience an upward trend while others experience a downward trend, then relatively many stores on a positive time trend will be selected into the second round. To analyze this, we run a pseudo-tournament among the stores in the control group. First, we group the control stores into poules in a similar way as the assignment of the treatment stores. We create 13 poules of 4 stores and two poules of 5 stores with similar average performance over the period August 2009 to August 2010. We identify for each of the poules the two stores with the highest performance during the first round of the experiment. Next, we compare the performance of the 'winners' and 'losers' of this pseudo-competition during the second round of the experiment with the performance of the real winners and losers from the treatment group. Figure 4 shows for each of these four groups the kernel densities of performance during the second round of the experiment. The performance distributions of the winners and losers of the pseudo-competition are very similar. Hence, in the control group, the stores that perform relatively well during the first round of the tournament do not show better or worse performance during the second round. Furthermore, the performance distribution of the first-round losers of the real tournament is similar to the performance distributions of the control stores. This again suggests that treatment stores not making it to the second round return to regular performance within two weeks of their elimination. In contrast, the second-round performance distribution of first-round treatment group winners is shifted to the right and has more mass between 2 and 2.4 as compared to the other groups. Hence, the second-round treatment effect is not caused by a selection of stores that experience a positive trend in performance.

Next, we analyze whether there are carry-over effects from the earlier experiment we did in this company. As described in section 5.3, all stores comprising the control group in the earlier experiment participated in the current tournament, as well as 29 randomly selected stores from the treatment group of the earlier experiment. Columns 1 and 2 of Table 3 show that the response of stores that did participate in the earlier experiment is not significantly different from the response of the other stores, neither in the first round nor in the second round. Hence, our current results are not affected by the earlier experiment.

Figure 5.4: Kernel density of performance during the second round

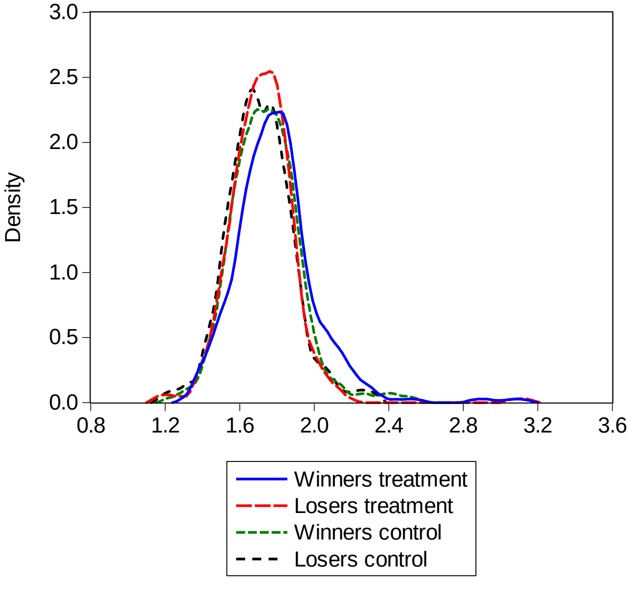

5.7.2 Noise

The first hypothesis of the model that we test is Proposition 1. Column 1 in Table 4 reports the results of estimating (5.4). We find that noisiness of performance measure APC has a negative effect on the response to the tournament. This negative effect is close to being significantly different from zero at conventional levels (p-value = 0.11). An increase in the level of noise by one standard deviation reduces the treatment effect by 1 percentage point. As we have taken up the variable noise in deviation from its mean, the first coefficient in column 1 gives the estimated treatment effect at the mean level of noise. This effect is about 1.3 percent and statistically different from zero. Wald tests show that the estimated treatment effect is significantly different from zero for stores with a standard deviation in APC up to 0.15 (i.e. for 70 percent of the stores).

Table 5.3: Carry-over effects

	Dependent variable: APC	
	(1)	(2)
Treatment * (Round 1 + Winners * Round 2)	0.029	
	(0.013)**	
Treatment * (Round 1 + Winners * Round 2) * Participant earlier experiment	-0.006	
	(0.016)	
Treatment * Round 1		0.017
		(0.011)
Treatment * Round 1 * Participant earlier experiment		-0.016
		(0.014)
Winners * Round 2		0.043
		(0.021)**
Winners * Round 2 * Participant earlier experiment		0.020
		(0.028)
Losers * Round 2	-0.014	-0.005
	(0.015)	(0.015)
Store-fixed effects	yes	yes
Week-fixed effects	yes	yes
Store-week observations	12079	12079
Stores	186	186
R^2	0.4471	0.4474

Standard errors clustered at the store level in parentheses.
Winners and Losers refer to the outcome of the first round of the tournament for the treatment stores.
Participant earlier experiment is a dummy variable that takes value 1 for treatment stores assigned to the treatment group in an earlier experiment ran in February 2010.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

In column 2 of Table 4, the estimation of (5.4) is separated by tournament round. We find a small and statistically insignificant effect of noise in the first round. In the second round, however, noise has a strongly significant, negative effect on performance in the tournament. Wald tests show that the second-round treatment effect is statistically different from zero for stores with a standard deviation below 0.16.[12] Hence, we find support for Proposition 1, in particular in the second round of the tournament. Note also that higher noise leads to weaker performance during the experimental period, underlining the importance of controlling for this time-specific effect.

Recall that noise in the performance measure APC is a pre-existing store characteristic, not randomly assigned. Hence, it is possible that noise is partially caused by or correlated with other store-specific characteristics. When these other store-

[12] A quadratic specification of the effect of sd_i does not improve the estimation.

Table 5.4: Noise in performance and the treatment effect

	Dependent variable: APC	
	(1)	(2)
Treatment * (Round 1 + Winners * Round 2)	0.024	
	(0.010)**	
Treatment * (Round 1 + Winners * Round 2) * StDev	-0.321	
	(0.198)	
(Round 1 + (Winners + Control) * Round 2) * StDev	-0.318	
	(0.104)***	
Treatment * Round 1		0.011
		(0.010)
Treatment * Round 1 * StDev		-0.188
		(0.199)
Round 1 * StDev		-0.365
		(0.111)***
Winners * Round 2		0.041
		(0.017)**
Winners * Round 2 * StDev		-0.632
		(0.282)**
(Winners + Control) * Round 2 * StDev		-0.269
		(0.140)*
Losers * Round 2	-0.015	-0.006
	(0.014)	(0.014)
Store-fixed effects	yes	yes
Week-fixed effects	yes	yes
Store-week observations	12079	12079
Stores	186	186
R^2	0.4494	0.4497

Standard errors clustered at the store level in parentheses.
Winners and Losers refer to the outcome of the first round of the tournament for the treatment stores.
StDev is a store's standard deviation of APC over the period August 2009 to August 2010.
This variable is mean-centered.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

specific characteristics affect stores' response to the tournament incentives, the effect of noise found in Table 4 might be spurious or estimated with bias. Insofar as these store characteristics are unobservable (at least for us), we cannot rule out this possibility. However, for observable store characteristics, this can be assessed. First, we run an OLS regression of our measure of noise on all available store characteristics (regression output can be found in the Appendix, Table A1). The observable store characteristics explain about 25 percent of the variation in noise across stores. Most explanatory power comes from the average level of performance APC and regional variation. Next, we take the residuals from this cross-section regression of noise, and use these residuals instead of the standard deviation of APC in estimating (5.4). Hence, we use the variation in noise that is not correlated with observable store characteristics. The results of this estimation are presented in Columns 1 and 2 in Table 5. We find that the estimates of the effect of residual noise on the response to competition are very similar to the estimates when using our standard measure of noise. This rules out that the negative effect of noise on the response to competition is caused by one or more of the observable store characteristics.

5.7.3 Prize spread

To analyse the effects of prize spread, we estimate the effects of the two treatments separately, as given by (5.5). The first column of Table 6 shows the estimated treatment effects over both rounds for the low-spread and high-spread treatments separately. Both treatments have a similar effect on performance, of around 1.5 percent in magnitude. Both estimates are significantly different from zero with a p-value of about 0.05.

Column 2 of Table 6 differentiates these estimates by tournament rounds. This estimation allows us to test the hypotheses that follow from Proposition 2 to 5. First, we focus on comparing the low-spread and the high-spread treatment. Propositions 2 and 3 predict better second-round performance in the high-spread treatment and better first-round performance in the low-spread treatment, respectively. We find that first-round performance in the low-spread treatment is indeed 0.8 percentage point better than in the high-spread treatment, but the difference is not statistically

Table 5.5: Noise uncorrelated with observables

	Dependent variable: APC	
	(1)	(2)
Treatment * (Round 1 + Winners * Round 2)	0.022	
	(0.011)**	
Treatment * (Round 1 + Winners * Round 2) * Residual noise	-0.342	
	(0.212)	
(Round 1 + (Winners + Control) * Round 2) * Residual noise	-0.225	
	(0.168)	
Treatment * Round 1		0.009
		(0.010)
Treatment * Round 1 * Residual noise		-0.203
		(0.184)
Round 1 * Residual noise		-0.230
		(0.150)
Winners * Round 2		0.036
		(0.017)**
Winners * Round 2 * Residual noise		-0.800
		(0.409)*
(Winners + Control) * Round 2 * Residual noise		-0.219
		(0.211)
Losers * Round 2	-0.015	-0.006
	(0.014)	(0.014)
Store-fixed effects	yes	yes
Week-fixed effects	yes	yes
Store-week observations	12017	12017
Stores	185	185
R^2	0.4485	0.4489

Standard errors clustered at the store level in parentheses.
Winners and Losers refer to the outcome of the first round of the tournament for the treatment stores.
Residual noise refers to the residuals of the OLS regression of stores' standard deviation of APC on all observable store-characteristics, as presented in Table A1. This variable is mean-centered.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

significant. In the second round, the treatment effect is 1 percentage point higher in the high-spread treatment, but again the difference is not statistically significant. Hence, both effects have the sign as predicted by theory, but the effects are small.

Next, we compare first and second-round performance within a treatment. Proposition 4 predicts that in the low-spread treatment, the first-round treatment effect is above the second-round treatment effect. The estimation results in column 2 of Table 6 shows that we actually find higher second-round performance, although the 1 percentage point difference is not statistically significant. Proposition 5 predicts that in the high-spread treatment, first-round performance should be below second-round performance. This is clearly borne out in column 2 of Table 6, where the difference between first and second-round performance is estimated at 2.7 percent, which is significant at a p-value of 0.022.

Two extensions of the basic model presented in section 5.4 might explain why the second-round treatment effect is higher than the first-round treatment effect in the low-spread treatment, in contrast to the model's prediction. First, there may be a selection effect. If stores differ in responsiveness to competition, then relatively responsive stores are selected into the second round. If so, we should compare the first and second-round treatment effect of the stores that won the first round. However, the stochastic nature of performance implies that we cannot simply compare the first and second-round performance of the first-round winners. Given that a store won the first round, its expected value of the stochastic component in APC during the first round is positive, yielding an upward bias in the estimate of the first-round effect. Here, we can use the pseudo-competition we conducted in the control group, as described in section 5.7.1, to assess the magnitude of this bias, as follows. The pseudo-competition gives us winners and losers of a competition purely determined by the stochastic component, in the same period as the first round of the tournament. We can compare the difference in performance between the winners and losers of the pseudo-competition to the difference in performance between winners and losers of the first round of the actual tournament. If stores are homogeneous, the theory as described in section 5.4 predicts that, while on average winners (losers) in the real tournament perform better than the winners (losers) in the pseudo-competition, the

Table 5.6: Estimated treatment effects: prize spread

	Dependent variable: APC		
	(1)	(2)	(3)
Low spread * (Round 1 + Winners * Round 2)	0.030		
	(0.016)*		
High spread * (Round 1 + Winners * Round 2)	0.026		
	(0.013)**		
Low spread * Round 1		0.021	
		(0.015)	
High spread * Round 1		0.006	
		(0.010)	
Control * Pseudo-winners * Round 1			0.072
			(0.012)***
Low spread * Losers * Round 1			0.015
			(0.012)
Low spread * Winners * Round 1			0.094
			(0.022)***
High spread * Losers Round 1			0.011
			(0.013)
High spread * Winners Round 1			0.071
			(0.012)***
Low spread * Winners * Round 2		0.038	0.041
		(0.025)	(0.025)
High spread * Winners * Round 2		0.057	0.058
		(0.023)**	(0.023)**
Losers * Round 2	-0.014	-0.005	-0.007
	(0.015)	(0.015)	(0.015)
Store-fixed effects	yes	yes	yes
Week-fixed effects	yes	yes	yes
Store-week observations	12079	12079	12079
Stores	186	186	186
R^2	0.4471	0.4474	0.4495

Standard errors clustered at the store level in parentheses.
Winners and Losers refer to the outcome of the first round of the tournament for the treatment stores.
Control * Pseudo-winners refers to the stores in the control group that 'won' the pseudo-competition.
Reference category for first-round effects in Column 3 are the 'losers' of the pseudo-competition.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

difference between winners and losers is similar across the treatment and control groups. If, on the other hand, stores differ in responsiveness to competition, the difference between winners and losers should be larger in the real tournament than in the pseudo-tournament.

Column 3 in Table 6 examines whether stores are heterogeneous in responsiveness to competition. The first five coefficients give the estimated performance during the first round of the experiment for five groups of stores, all relative to the performance of the stores in the control group that 'lost' the pseudo-competition. First, the 'winners' of the pseudo-competition perform about 4 percentage points better than the 'losers'. Comparing the difference in performance between the first-round winners and losers in the treatment groups, we see that in the low-spread treatment this difference is marginally higher at 4.3 percentage points, while in the high-spread treatment it is even smaller at 3.3 percentage points. These differences are nowhere close to being statistically significant. Hence, we find no evidence for differences in responsiveness across stores.

An alternative explanation is that winning a competition may provide employees with non-monetary benefits such as status, social recognition, or simply the joy of winning (Auriol and Renault 2008, Besley and Ghatak 2008, Frey and Neckermann 2008, Moldovanu et al. 2007). Several recent studies suggest that these non-monetary benefits are substantial, by showing that people respond to competition even when there is no money at stake (i.e. when only relative performance information is provided), see Azmat and Iriberri (2010), Blanes i Vidal and Nossol (2009), Delfgaauw et al. (2009), Kosfeld and Neckermann (2010), and Sheremeta (2010). Accepting the presence of non-monetary utility of winning a competition, the result that the first-round treatment effect is above the second-round treatment effect in the low-spread treatment would suggest that winning the second round yields higher non-monetary utility than winning the first round. Note that the addition of a non-monetary benefit of winning the second round of the tournament to the basic model in section 5.4 does not change the predictions of the effects of changes in the prize spread. In particular, the difference between first-round and second-round performance in the high-spread tournament should be larger than this difference in the low-spread tournament. Computed from the estimates in the second column of Table 6, the magnitude of this difference-in-differences is about 1.8 percentage points, but it is not statistically significant.

5.8 Concluding remarks

Examining whether workers respond as predicted to tournament incentives in their natural working environment is important for linking tournament theory to organizational policies regarding wages and promotions. We have designed a natural field experiment in a private company to test several predictions on the effects of prize structure and noise in an elimination tournament. As predicted, we find that increased convexity of the prize spread increases second-round performance at the expense of first-round performance. Furthermore, workers with relatively volatile performance hardly respond to tournament incentives, while workers whose performance measure is stable increase performance significantly. One finding is at odds with standard theory: In a treatment with equal prizes for winning the first and second round, performance is better in the second round than in the first round. This suggests that workers attach non-monetary benefits to becoming the overall winner of the tournament.

The magnitude of the effects we find is not particularly large. The average treatment effect on the performance measure APC (Average number of Products Sold per Client) is about 1.5 percent. In the end, the company's management cares about sales revenues rather than APC. However, we find no effect of the experiment on sales. Apparently, workers have means to increase APC without increasing revenue, suggesting that APC is prone to gaming.

5.A Appendix

Table 5.7: OLS of noise on observables

	Dependent variable: standard deviation of APC
	(1)
Mean APC	0.183
	(0.044)***
Number of employees	0.001
	(0.003)
Number of employees in full-time equivalents	0.000
	(0.000)
Average age employees	-0.001
	(0.002)
Percentage of female employees	0.010
	(0.017)
Average tenure of employees	0.005
	(0.003)*
Brand 2	-0.033
	(0.018)*
Constant	-0.165
	(0.083)**
Regional dummies	yes
Stores	185
R^2	0.2453

Standard errors in parentheses.
Mean standard deviation of APC are based on the period August 2009 to August 2010.
The personnel variables are extracted from the companies' database as of September 1, 2010.
The personnel information is missing for 1 store in the analysis.
Brand 2 is a dummy variable for stores operating under the companies' smaller brand name.
***, **, * denote statistically significant effects at the 1%, 5%, and 10% level, respectively.

Chapter 6

Summary and Directions for Further Research

A central problem in economics is the problem of agency. A principal hires an agent who is supposed to take care of the principal's interests, but the agent's interests differ to some extent from the principal's. Economists have generally considered the employer-employee relation as a typical example of an agency relationship, and stressed the importance of financial incentives to align the interests of both parties. Although empirical studies show that pay-for-performance indeed motivates workers to exert effort, more recent papers acknowledge that workers are not only driven by their own material self-interest. Several other, often immaterial, considerations also play an important role, such as sensitivity to social norms (Sliwka 2007), reciprocity (Falk and Fischbacher 2006), status concerns (Moldovanu 2007), and intrinsic motivation (Delfgaauw and Dur 2007). Another recent development is the increased use of field experiments to empirically investigate the effects of incentives. The advantage of experiments over naturally occuring data is that they greatly facilitate causal inference.

This thesis contributes to this emerging literature by studying how employers motivate employees to exert effort on the job. The first part of this thesis, chapters 2 and 3, analyzes agency problems when workers are reciprocal towards the employer's kindness. The second part of this thesis, chapters 4 and 5, presents two field experiments designed to test some basic predictions of tournament theory. In this chapter, I summarise the main findings and provide some avenues for further research.

6.1 Summary

The key assumption in the first part of this thesis is that workers are reciprocal. Reciprocity usually means that people respond to kind acts by being kind in return, and likewise for unkind acts. However, it is not always obvious what constitutes a 'kind act', because the intentions behind an action matter just as well as the distributional outcomes, see e.g. Charness and Levine (2007). Thus, it seems fair to assume that people do not simply reciprocate favors, but instead care about whether a person is a kind person. We refer to this type of reciprocity as conditional altruism, defined as the willingness to promote the welfare of a kind person and to reduce the welfare of an unkind person.

I use this definition of reciprocity in chapter 2 to investigate the relation between reciprocity and incentive pay, while allowing for worker heterogeneity. Although there is a sizeable literature on how employers can trigger effort by paying high wages, the literature pays little attention to the question how wages and explicit incentives interact when workers are reciprocal. I address this gap in the literature by analyzing a formal principal-agent model, assuming that workers are risk-averse and conditionally altruistic. Conditionally altruistic workers care more for their employer when they are convinced that the employer cares for them. I assume that some, but not all, workers are conditionally altruistic, because a typical finding in laboratory experiments is that not all individuals are equally motivated by reciprocal tendencies. The employer can be of two types, namely altruistic or egoistic. He offers a contract consisting of a base salary and a share of the production value. He is hiring for a single position from a large pool of workers. Workers interpret the employer's contract offer as a signal of his kindness.

I find that employers signal their altruism by offering relatively weak incentives and a relatively high base salary. Although a contract combining steep incentives with an appropriate base salary may just as well signal the employer's care, offering relatively weak incentives is optimal because conditionally altruistic workers do not need strong incentives when they sense that their employer cares for them. The level of the total expected compensation offered positively depends on the additional effort induced by reciprocity, as credible signalling rules out that an altruistic principal makes more profits. This implies, however, that altruistic employers may pay less in expectation than egoistic employers: because altruistic employers offer weaker incentives, workers may exert lower effort. In that case the altruist's gift consists of weak incentives, implying little pressure to exert effort and low exposure to risk.

I also analyze employer's incentives to screen workers. Since some workers do not reciprocate the employer's altruism, the employer may find it optimal to write a contract that simultaneously signals his altruism and screens conditionally altruistic worker types. I show that such a contract is characterised by excessively strong incentives, a relatively high base salary, and a relatively high expected total compensation. As conditionally altruistic workers exert more effort than selfish workers when they are convinced that their employer cares for them, accepting a contract specifying strong incentives is relatively more attractive for conditional altruists than selfish workers. The paradox is that strong incentives are offered to attract the employees who need them the least.

While in chapter 2 I depart from the assumption that individuals are only concerned about their narrow self-interest, I maintained the traditional economic view of the employer as a passive contract writer. The employer shows his care for the worker by offering a generous contract, but has no other means of doing so. Although a standard assumption in the economic literature on reciprocity, in practice employers have several non-monetary tools of management at their disposal, which I refer to as management attention. In chapter 3, my co-authors and I study optimal incentive contracts for workers who are reciprocal to management attention. In contrast to chapter 2, we abstract from possible intentions behind an action, instead we assume that management attention always raises worker's well-being and is reciprocated with higher effort. As in chapter 2, we assume that the manager is also the owner of the firm, implying that he is concerned about the firm's profits.

We find that when neither worker's effort nor manager's attention can be contracted, a double moral-hazard problem arises. Workers lack the incentives to provide the first-best level of effort, whereas managers have insufficient incentives to provide attention. Paying workers according to their individual performance provides them with an incentive to exert more effort, but on the other hand reduces the manager's incentive to provide attention, as there is less at stake for the manager. In a multi-agent setting, the manager can resolve the problem by committing to promotion incentives. Thus, we would expect that more reciprocal workers are relatively unlikely to receive bonus incentives, but more likely to receive promotion incentives. We empirically examine these predictions using data from the German Socio-Economic Panel. This dataset contains a direct measure of an individual's reciprocity, in addition to information on his remuneration. We find that workers who are more reciprocal are significantly more likely to receive promotion incentives, while there is no significant relation for individual bonus pay.

The second part of this thesis, chapters 4 and 5, differs from the first part by the topic we examine and the method used. Instead of theoretically investigating the consequences of reciprocity for economic outcomes, we conduct two natural field experiments in a Dutch retail chain to examine how tournaments affect performance. In chapter 4 we study how intermediate performance feedback affects subsequent performance. For instance, a big prize for winning a tournament is unlikely to motivate contestants who learn during the contest that they have no chance of winning anymore. Likewise, the prize is unlikely to motivate contestants who learn that they are comfortably ahead. These predictions generalize to other non-linear pay-for-performance plans. To test these predictions, we conduct a field experiment among 189 stores of a Dutch retail chain. Testing these predictions is more complicated than it may seem at first sight. As a store's sales figures are serially correlated, the estimates of the effect of intermediate relative performance on subsequent performance are biased. We address this issue by introducing a relative performance pay scheme in a randomly selected subset of the retail chain's stores. Specifically, employees in the treatment stores could win a bonus by sufficiently outperforming three comparable stores from the control group over the course of four weeks. Treatment stores received weekly feedback on relative performance, while control stores were kept unaware of their involvement. As a result, the performance of the control stores generates exogenous variation in relative performance.

Using the performance of the control stores as an instrument for the effect of intermediate relative performance, we confirm the theoretical predictions. We find that treatment stores that lag far behind do not respond to the bonus, while the responsiveness of treatment stores close to winning a bonus increases in relative performance. Specifically, we find that one percentage point increase in relative performance increases subsequent performance by 0,72%. We cannot test, however, whether treatment stores well ahead of their closest competitor reduce their efforts, as none of the treatment stores happened to be in this comfortable position. On average, the introduction of the relative performance pay scheme does not lead to higher performance, as only few treatment stores were in a position to win a bonus.

In the field experiment presented in chapter 5, our main interest is to identify how the distribution of total prize money over the two rounds of an elimination tournament influences performance over the course of the tournament. Standard tournament theory predicts that a convex prize structure (i.e. a higher prize in the second round than in the first round) leads to better second-round performance at the expense of first-round performance. A second objective is to investigate whether the incentive effect of the tournament depends on noise in the performance measure. Theory predicts that noise dilutes incentives to perform, as noise reduces the probability that additional effort leads to a win.

To test these hypotheses, we run an elimination tournament among a randomly chosen subset of the company's 208 stores. The tournament consists of two rounds, both lasting four weeks. In the first round, the 144 participating stores are assigned to poules of four stores that are comparable in terms of historical performance. After the first round, the bottom-two stores of each poule are eliminated, whereas the top-two stores of each poule win a prize and continue to the second round. In the second round, stores are once more grouped into poules of four comparable stores. The two best-performing stores of each poule again win a prize. Participating stores are assigned to two different treatments that differ by the prize spread only. We keep the total amount of prize money constant. In the low-spread treatment, prizes are identical in the two rounds, whereas in the high-spread treatment prizes are increasing over the two rounds. To test whether noise negatively affects the stores' response to the tournament, we divide the treatment stores in a high-noise and low-noise group before assigning them to a treatment condition. We use the variance in performance prior to the experimental period as our measure of noise. We keep both noise groups separated throughout the tournament, implying that all stores in a given poule are exposed to a similar amount of noise in the performance measure.

We find that second-round performance is around 1 percentage point higher in the high-spread treatment as compared to the low-spread treatment, while first-round performance is 0.8 percentage point lower. However, these differences are not statistically significant. The average incentive effect of the tournament is approximately 1,5%, irrespective of the prize spread. Noise in the performance measure negatively influences the treatment effect, potentially offsetting it. However, this happened in only 30% of the cases, implying that the tournament significantly improved performance of the large majority of participants.

6.2 Directions for further research

In chapters 2 to 3 of this thesis, we have studied how managers can use reciprocity to motivate workers. We thus concentrated on positive reciprocity: people's tendency to reward favors. However, empirical evidence shows that negative reciprocity, the inclination to reciprocate unfriendly behavior, is a much stronger force (e.g. Kube et al. 2006). A logical extension would thus be to examine how managers can prevent negative reciprocity when taking unpopular measures is required.

Although studying negative reciprocity seems a promising direction for research, several interesting issues remain when we concentrate on positive reciprocity. An important question is 'the currency of reciprocity' (Kube et al., 2008): how can employers most efficiently signal to their employees that they care about worker's welfare? Instead of money or personal attention, employers may consider giving in-kind gifts (Kube et al., 2008) or granting more decision rights to workers. Thus, relevant questions demanding empirical and theoretical investigation are how employers mix these instruments, why certain combinations are more effective than others, and under which conditions. Another avenue for further research is to recognize that not only workers are reciprocal, but that managers have reciprocal feelings as well. This may have serious consequences for the effort managers exert on the job, the profitability of the organisation, and whether managers should have discretion over pay or not.

It may also be worthwhile to analyse an important aspect of management attention in more detail, namely praise for a job well-done. Although praise is often considered an important motivator in the practioner's literature, the literature in economics is largely silent on the issue. It would be interesting to develop micro-foundations why praise may work as a motivational tool.

Finally, although the first part of my thesis is devoted to analyzing the implications of reciprocity in employer-employee relations, reciprocity may also play a crucial role in other relations. For instance, reciprocity is considered to be an important factor in relations between colleagues (Dur and Sol, 2010), or in relations between a firm and its suppliers or its customers (Huck and Tyran 2007). Several highly relevant questions in this area are worth investigating in more detail. For instance, how do good personal relations between colleagues affect organisational structure and performance? How does reciprocity affect the efficient functioning of markets, in particular in situations of asymetric information?

In the second part of my thesis, I have presented two field experiments testing predictions derived from tournament theory. Obviously, the experimental method can also be used to test other predictions, in particular those derived from non-standard economic theory. For instance, what are the consequences of incentives for social relations in the workplace? Do organisations offering relatively steep incentives attract workers with a different personality? Do in-kind benefits have a larger impact on job motivation than monetary rewards, as found by Kube et al. (2008), and why? And do the effects persist over time? These topics are left for further research.

Samenvatting (Summary in Dutch)

Introductie

Traditioneel hebben economen grote nadruk gelegd op mogelijk belangenconflict tussen werknemer en werkgever. Werknemers worden geacht het belang van hun werkgever te bevorderen, maar dit kan strijdig zijn met hun eigen voorkeuren en interesses. Ter illustratie, werknemers zullen wellicht vaker grappen tappen bij de koffieautomaat dan hun baas lief is. Om dit te voorkomen zou de werkgever de koffieautomaat dicht bij zijn eigen kantoor neer kunnen zetten, maar er zijn diverse situaties te bedenken waarin de werkgever de werknemers niet zo eenvoudig kan controleren. De oplossing die economen normaliter aandragen is daarom een andere: het belangenconflict kan worden opgeheven door de werknemer hetzelfde financiële belang bij goede prestaties te geven als de werkgever heeft. Met andere woorden, de beloning van de werknemer moet gebaseerd zijn op zijn prestaties, zodat hij een financiële prikkel heeft om het door de werkgever gewenste gedrag te vertonen. Het belang van financiële prikkels is aangetoond in diverse studies (zie Prendergast 1999 en Lazear en Oyer 2009 voor een overzicht). Een indrukwekkend voorbeeld wordt gegeven door de vroege econoom Edwin Chadwick (1800-1890). Hij slaagde erin om het percentage Britse criminelen dat het transport naar Australië overleefde te verhogen van 40% naar 98%, simpelweg door kapiteins te betalen voor iedere veroordeelde die Australië levend bereikte, in plaats van voor iedere veroordeelde die in Engeland aan boord ging. Meer recent zijn de studies van Lazear (2000) en Shearer (2004), die beide kijken naar de effecten van het betalen van stukloon in plaats van vast loon. Beide studies vinden een productiviteitsstijging van ongeveer 20% als gevolg van de introductie van stukloon.

Echter, weinig economen zullen beweren dat bedrijven alleen op monetaire prikkels vertrouwen om werknemers te motiveren. De grote meerderheid van de werknemers heeft er nauwelijks financieel belang bij om zich in te spannen. In Westerse landen krijgt minder dan de helft van alle werknemers een salaris dat expliciet gerelateerd is aan zijn prestaties, en wanneer dat wel het geval is, dan is de hoeveelheid geld die op het spel staat meestal gering, namelijk minder dan 5% van het vaste salaris, zie Piekkola (2005), Lemieux et al. (2007) en Belfield en Marsden (2003). Bovendien vinden noch werknemer noch werkgever financiële prikkels de belangrijkste motivator, zie de enquêtes van Minkler (2004) en Campbell en Kamlani (2007).

Gezien deze observaties is het niet verrassend dat economen het beeld van werknemers en werkgevers als respectievelijk onbetrouwbare luiaards en passieve contractschrijvers recentelijk bijgesteld hebben. Economen erkennen dat werknemers niet uitsluitend geïnteresseerd zijn in het maximaliseren van hun eigen materiële welvaart, en dat diverse andere motieven een belangrijke rol spelen op de werkplek. Ter illustratie, recente papers houden er rekening mee dat werknemers intrinsiek gemotiveerd kunnen zijn (Delfgaauw en Dur 2007), dat de mate waarin werknemers geneigd zijn zich in te spannen mede afhangt van sociale normen (Sliwka 2007), en dat werknemers behoefte hebben aan respect en erkenning van hun werkgever (Ellingsen en Johannesson 2007). Al deze niet-financiële overwegingen kunnen soms een verrassende samenhang vertonen met financiële prikkels. Een bekend voorbeeld is dat financiële prikkels soms ten koste gaan van de intrinsieke motivatie voor een bepaalde taak (Frey and Jegen, 2001), of opgevat worden als signaal dat zelfzuchtig gedrag blijkbaar de heersende norm is (Sliwka 2007).

In het licht van deze ontwikkelingen is het wellicht niet verwonderlijk dat economen ook hun toevlucht hebben genomen tot nieuwe methoden van onderzoek. Economen maken tegenwoordig steeds meer gebruik van de experimentele methode om de verschillende motieven die bij het menselijk gedrag een rol spelen te kunnen onderscheiden. Een belangrijk voordeel van de experimentele methode is dat de onderzoeker de omgeving zodanig kan manipuleren dat het mogelijk is om causale verbanden vast te stellen. Met bestaande datasets is dit vaak niet mogelijk vanwege verstorende omgevingsfactoren. Als we bijvoorbeeld het effect van prestatieloon op het bedrijfsresultaat vast willen stellen, is het niet voldoende om bedrijven met en zonder prestatieloon te vergelijken. De reden is dat de verschillen in de manier waarop bedrijven hun personeel belonen, niet toevallig zijn: bedrijven die prestatieloon invoeren zijn wellicht bedrijven met een meer prestatiegerichte bedrijfscultuur, of er is bijvoorbeeld alleen draagvlak voor het invoeren van prestatieloon als het bedrijf goed presteert. Het simpelweg vergelijken van bedrijven met en zonder prestatieloon zou dus tot onjuiste conclusies over de effecten ervan kunnen leiden. Om dezelfde reden is het problematisch om de resultaten van een bedrijf voor en na de invoering van prestatieloon met elkaar te vergelijken. Een verandering in de prestaties zou namelijk ook veroorzaakt kunnen worden door een toevallige verandering in de vraag naar het product, of een andere beleidswijziging die gelijktijdig wordt ingevoerd. Om een causaal verband vast te stellen is het altijd nodig te kunnen bepalen wat er gebeurd zou zijn als een bepaalde verandering, zoals de invoering van prestatieloon, niet had plaatsgevonden. Experimenten maken dit eenvoudig doordat de onderzoeker willekeurig sommigen wel en anderen niet aan een bepaalde verandering blootstelt. Zolang beide groepen elkaar niet beïnvloeden, is een zuivere schatting van de effecten mogelijk. Om deze reden nemen economen steeds vaker hun toevlucht tot laboratorium of veld-experimenten.

Dit proefschrift draagt op diverse manieren bij aan de zojuist beschreven literatuur. In het eerste deel van het proefschrift, de hoofdstukken 2 en 3, analyseren mijn co-auteurs en ik het bovengenoemde belangenconflict tussen werkgever en werknemer wanneer werknemers reciproque zijn. Daarmee bedoel ik dat werknemers zich coöperatiever opstellen als ze merken dat hun baas om hen geeft. Het doel van de analyse is om te bepalen hoe deze neiging om de ander met gelijke munt terug te betalen de optimale beloningsstructuur beïnvloedt. In het tweede deel van dit proefschrift, de hoofdstukken 4 en 5, beschrijf ik twee veldexperimenten die gehouden zijn in een Nederlandse winkelketen. Beide veldexperimenten testen specifieke theoretische voorspellingen over de effecten van verkoopwedstrijden. Ik zal nu kort de literatuur over reciprociteit beschrijven. Vervolgens zal ik een overzicht geven van de inhoud van ieder hoofdstuk.

Reciprociteit

De sterke nadruk van economen op de financiële drijfveren van werknemers heeft ertoe geleid dat economen lange tijd weinig aandacht besteed hebben aan andere belangrijke motieven die het menselijk gedrag kunnen verklaren. Zoals hierboven beschreven, heeft de recente literatuur hier meer oog voor. Een van deze drijfveren die recentelijk veel onderzocht is, is reciprociteit. Reciprociteit ofwel wederkerigheid betekent dat mensen de neiging hebben anderen te behandelen zoals ze zelf door hen behandeld worden. De neiging om de welvaart van een vriendelijk persoon te bevorderen wordt ook wel positieve reciprociteit genoemd, terwijl de neiging om de welvaart van een onvriendelijk persoon te verminderen als negatieve reciprociteit bekend staat (Dohmen et al. 2009). Beide soorten reciprociteit vinden we terug in bekende spreekwoorden als "wie goed doet, goed ontmoet", en "wie kaatst, moet de bal verwachten".

Het belang van reciprociteit op de werkplek is met name onder de aandacht van economen gebracht door een artikel van George Akerlof in 1982. In dit artikel beargumenteert hij dat de inspanningen van werknemers en de betrekkelijk goede behandeling van werknemers door de werkgever gezien kunnen worden als wederzijdse geschenken. Werknemers produceren meer dan strict noodzakelijk, terwijl werkgevers hun werknemers beter behandelen dan vereist. Arbeidscontracten kunnen dus beschouwd worden als een uitruil van geschenken, waarbij wederkerigheid de norm is.

Dit concept van wederzijdse vriendelijkheid dateert van ver voordat Akerlof deze zienswijze in de economie introduceerde. Homans, een sociaal psycholoog, beschreef al in 1958 sociaal gedrag als een uitruil van materiële en immateriële goederen. Binnen de sociologie maakt Blau (1964) onderscheid tussen economische en sociale ruil. Sociale ruil heeft als kenmerk dat de wederzijdse verplichtingen niet gespecificeerd zijn. Hoewel er een algemene verwachting is dat wanneer de een de ander een gunst verleent, de ander hier iets tegenover zal stellen, is het niet van tevoren vastgelegd wat die ander terug zal doen en wanneer. Hier kan ook redelijkerwijs niet over onderhandeld worden. Volgens Blau (1964) proberen mensen het evenwicht te bewaren tussen wat ze in een relatie stoppen en wat ze eruit halen. Met andere woorden, het ontvangen van een gunst schept een morele verplichting iets terug te doen. Het krijgen van een gift is dan ook een twijfelachtige eer, wat geïllustreerd wordt door de etymology van het woord ‘gift’. In oud-Germaans en Grieks, wordt hetzelfde woord gebruikt voor gift en vergif.

Economen zijn over het algemeen van mening dat reciprociteit van groot belang is voor het begrijpen van arbeidscontracten. Vele laboratorium experimenten hebben Akerlof’s hypothese bevestigd, namelijk dat werknemers bereid zijn om zich meer in te spannen wanneer ze een hoger loon ontvangen (Fehr en Gächter, 1998).[1] Bovendien, een veldexperiment van Kube et al. (2008) laat zien dat een gift in natura leidt tot een aanzienlijke stijging van de productiviteit van werknemers, meer nog dan een salarisverhoging van dezelfde waarde. Verder is in diverse laboratorium-experimenten aangetoond dat een genereus contract werknemers sterker motiveert dan contracten die boetes in het vooruitzicht stellen wanneer de inspanningen van de werknemers achterblijven bij een van tevoren afgesproken hoeveelheid. Tenslotte laten enquêteresultaten zien dat mensen die sterker geneigd zijn goed voor goed te vergelden meer verdienen, een kleinere kans hebben werkloos te raken, en meer overwerken, met name als ze hun loon als eerlijk beschouwen (Dohmen et al. 2009). Deze bevindingen suggereren dat werkgevers van reciprociteit kunnen profiteren, zolang ze hun personeel goed behandelen tenminste. Aan de andere kant doen werkgevers er goed aan om geen negatief reciproque gevoelens op te wekken. Het is om deze reden dat leidinggevenden zeer terughoudend zijn in het verlagen van lonen (Bewley 1999). Kube et al. (2008) vinden in een veldexperiment dat werknemers inderdaad sterk negatief reageren op een loon dat achterblijft bij de gewekte veranderingen. Ook is er overvloedig bewijs dat mensen oneerlijke voorstellen in onderhandelingsspellen verwerpen, ook als ze zichzelf daarmee benadelen (Camerer en Thaler 1995).

Reciprociteit hangt echter niet alleen af van de uitkomsten van bepaalde acties, zoals gesuggereerd door bovenstaande literatuur. De intentie van een bepaalde handeling is eveneens van belang, zoals aangetoond in diverse experimenten, bijvoorbeeld Brandts en Charness (2003), Falk et al. (2003 en 2008), en Charness en Levine

[1] Veldexperimenten verlenen echter minder bewijs voor Akerlof’s hypothese, zie bijvoorbeeld Gneezy en List (2006) en List (2006).

(2007)). Economen hebben daarom modellen ontwikkeld die rekening houden met het belang van de intenties van individuen. Reciprociteit is in dat geval definieerd als de gedragsreactie op een actie die als vriendelijk of onvriendelijk wordt ervaren (Falk en Fischbacher 2006). Of een actie als vriendelijk of onvriendelijk wordt gezien hangt zowel af van de gevolgen als van de intentie van de actie, zie bijvoorbeeld Rabin (1993) en Falk en Fischbacher (2006). Deze modellen zijn echter zeer complex. Een eenvoudigere benadering is die van Levine (1998). In het model van Levine gaat het er niet zozeer om of mensen zich vriendelijk of onvriendelijk gedragen, maar of ze feitelijk om de ander geven. Met andere woorden, Levine veronderstelt dat mensen conditioneel altruistisch zijn: ze geven meer om het welzijn van de ander, als ze merken dat de ander ook om hen geeft. Vriendelijk gedrag wordt dus niet automatisch overeenkomstig beantwoord, maar alleen wanneer dit een overtuigend signaal is van de altruistische gevoelens van de ander. Deze benadering van reciprociteit komt overeen met de bevindingen binnen de organisatiepsychologie: als werknemers merken dat hun leidinggevende of de organisatie geeft om hun welzijn, dan leidt dat tot betere prestaties en een grotere loyaliteit (Rhoades en Eisenberger 2002 en Cropanzano en Mitchell 2005).

Overzicht van de hoofdstukken

Hoofdstuk 2: reciprociteit en prestatieloon

Het eerste deel van dit proefschrift draagt op diverse manieren bij aan de literatuur over reciprociteit in organisaties. In hoofdstuk 2 stel ik een theoretisch model op om de relatie tussen reciprociteit en prestatieloon te onderzoeken. Wanneer mensen reciproque zijn wordt het betalen van een riant vast salaris veelal gezien als een waardig substituut voor prestatieloon, maar er is natuurlijk geen reden waarom werkgevers niet beide instrumenten zouden gebruiken om werknemers te motiveren. De weinige bestaande literatuur levert geen eenduidige conclusies over wat het best mogelijke beloningsbeleid is. Bellemare en Shearer (2011) concluderen dat werkgevers beter reciprociteit kunnen opwekken door de prestatieafhankelijke component te verhogen dan het vaste salaris. Englmaier en Leider (2008) vinden het tegenovergestelde. Ten opzichte van de bestaande literatuur kent mijn model twee belangrijke vernieuwingen. Ten eerste houdt mijn model rekening met het belang van intenties. Daarnaast veronderstel ik dat niet alle werknemers reciproque zijn. Dit reflecteert de bevindingen uit diverse lab-experimenten (bijvoorbeeld Fischbacher et al. 2001): sommigen mensen gedragen zich egoistisch onafhankelijk van het gedrag van anderen.

Om intenties te modelleren gebruik ik de benadering van Levine (1998). Ik veronderstel dat de werkgever egoistisch of altruistisch is, en dat een bepaalde fractie van werknemers altruistisch is als ze merken dat de werkgever om hen geeft. Met andere woorden, sommige, maar niet alle, werknemers zijn conditioneel altruistisch. De werkgever huurt 1 werknemer uit een grote groep potentiële werknemers. Hij biedt hen een arbeidscontract aan dat bestaat uit een vast salaris en een prestatie-afhankelijk deel. Werknemers bepalen aan de hand van het arbeidscontract of de werkgever om hen geeft of niet, en degenen die daar gevoelig voor zijn reageren hierop door zich meer in te spannen. Ik onderzoek hoe altruistische werkgevers zich het beste kunnen onderscheiden van egoisten. De optimale manier om dit te doen, is een relatief hoog basissalaris aanbieden in combinatie met een geringe prestatie-afhankelijke component. De reden is dat het relatief hoge basissalaris werknemers overtuigt van de goede bedoelingen van de werkgever, zodat sterke financiële prikkels overbodig zijn. Daarnaast onderzoek ik of de werkgever een beloningsstructuur kan bedenken die laat zien dat hij om zijn personeel geeft, maar tegelijkertijd zelf-selectie bewerkstelligt van werknemers die gevoelig zijn voor het altruisme van de werkgever. Beide doelen kunnen verwezenlijkt worden door een betrekkelijk hoog basissalaris te combineren met een grote prestatie-afhankelijke component.

Hoofdstuk 3: management attentie en prestatieloon

In hoofdstuk 2 wordt een al dan niet goede behandeling van werknemers gedefinieerd in monetaire termen. Hoewel dat binnen de economische literatuur een gebruikelijke veronderstelling is, is reciprociteit natuurlijk niet beperkt tot de financiële aspecten van een baan. Stelt u zich bijvoorbeeld een werkgever voor die een bovengemiddeld hoog loon betaalt, maar nauwelijks aandacht geeft aan zijn personeel, en als hij dat dan eens doet, dan is dat alleen om op barse toon te vertellen dat iemand een fout gemaakt heeft. Het is moeilijk voor te stellen dat werknemers bereid zijn om voor deze werkgever een stapje extra te doen: het tegenovergestelde ligt meer voor de hand. Om deze reden erkennen studies in organisatiepsychologie het belang van 'zachte', sociale aspecten voor de arbeidsmotivatie, zie bijvoorbeeld Cropanzano en Mitchell (2005).

Geïnspireerd door voornoemde literatuur veronderstel ik -samen met mijn co-auteurs- in hoofdstuk 3 dat werknemers positief reageren op de aandacht die ze van hun werkgever ontvangen. Met aandacht bedoelen we diverse niet-financiële aspecten van management, zoals het geven van waardering, erkenning voor een goede prestatie, persoonlijke aandacht etc. We veronderstellen dat aandacht werknemers gelukkiger maakt, en dat werknemers op aandacht reageren door zich sterker in te spannen. Vervolgens analyseren we de prikkels voor de werkgever om aandacht te geven en de implicaties voor de optimale beloningsstructuur.

De belangrijkste bevinding is dat individueel prestatieloon ten koste gaat van de prikkels voor de werkgever om aandacht te geven. Naarmate de werknemer een groter deel van zijn productiewaarde mag houden, staat er voor de werkgever zelf financiëel minder op het spel, wat impliceert dat hij minder reden heeft om de werknemer middels persoonlijke aandacht aan te zetten tot betere prestaties. Wanneer er meerdere werknemers zijn, kan de werkgever dit oplossen door een promotie in het vooruitzicht te stellen. We verwachten dan ook dat naarmate werknemers een sterkere reciproque neiging hebben, het waarschijnlijker is dat ze promotie kunnen maken, en minder waarschijnlijk dat individueel prestatieloon onderdeel uitmaakt van hun beloning. Wij testen deze voorspelling met data van de German Socio-Economic Panel, omdat deze dataset een maatstaf heeft voor reciprociteit en informatie over iemands beloningspakket. We vinden inderdaad dat naarmate een individu sterker reciproque is, het waarschijnlijker is dat hij promotie kan maken. We vinden echter geen duidelijke relatie tussen iemands reciprociteit en de kans dat die persoon prestatieloon ontvangt.

Hoofdstuk 4: tussentijdse feedback

Het tweede deel van dit proefschrift bestudeert het effect van verkoopwedstrijden door middel van twee veldexperimenten. Anders dan het eerste deel van het proefschrift is dit onderzoek dus empirisch van aard, en niet theoretisch. In hoofdstuk 4 bekijken we hoe tussentijdse feedback over de relatieve prestaties de daaropvolgende prestaties beïnvloedt. Om dit te onderzoeken, organiseren we een verkoopwedstrijd tussen soortgelijke filialen van een winkelketen, waarbij deelnemende filialen wekelijks tussenstanden ontvangen. De theorie voorspelt dat deelnemers die zien dat ze geen kans meer hebben om te winnen, afhaken, terwijl deelnemers die leren dat ze ver voorstaan op hun lauweren zullen gaan rusten. Het aantal eerdere studies over dit onderwerp binnen de context van een bestaande organisatie is zeer gering, namelijk Casas-Arce en Martinez Jerez (2009) en Frank en Obloj (2011). Ons onderzoek onderscheidt zich van deze studies door de gevolgde methode. Anders dan genoemde studies gebruiken wij de experimentele methode, waarbij de aanwezigheid van een controlegroep een betrouwbare schatting van het effect van verkoopwedstrijden op de omzet mogelijk maakt. Verder is het experiment zodanig opgezet dat we eenvoudig kunnen corrigeren voor de aanwezigheid van seriële correlatie. Seriële correlatie betekent dat toevallige schommelingen in de omzet over tijd met elkaar samenhangen: een tegenvallende omzet in week 1 betekent dat de omzet in week 2 waarschijnlijk ook tegenvalt (of juist meevalt, afhankelijk van de richting van de correlatie). De moeilijkheid is om de bestendigheid van toevalsfactoren te onderscheiden van het effect van de tussenstanden. Wanneer een aantal winkels twee weken achtereenvolgens slecht presteert, is dit het gevolg van toeval, of heeft de lage positie in de tussenstanden in de eerste week het winkelpersoneel ontmoedigt, en derhalve een slechte prestaties in de tweede week veroorzaakt?

Het experiment werd gedaan in een Nederlandse winkelketen die muziek, computerspellen en films verkoopt. Een willekeurig gekozen groep winkels kreeg de mogelijkheid om in een verkoopwedstrijd een bonus te verdienen (treatmentgroep), de overige winkels functioneerden als controlegroep. Om de bonus te verdienen moesten winkels uit de treatmentgroep drie vergelijkbare winkels uit de controlegroep verslaan. Als prestatiemaatstaf diende de omzet ten opzichte van het salestarget gedurende de vier weken van het toernooi. De winkels uit de treatmentgroep kregen gedurende het toernooi iedere week een poster met tussenstanden. De winkels in de controlegroep waren niet op de hoogte van hun betrokkenheid bij het toernooi. Er is dus sprake van een semi-toernooi: winkels uit de treatmentgroep nemen het op tegen winkels die zelf niet meedoen. De reden voor deze opzet is dat de prestaties van de winkels uit de controlegroep niet beinvloedt worden door de prestaties van de winkels uit de treatmentgroep. We kunnen nu testen of dat andersom (via de communicatie van de tussenstanden) wel het geval is, met andere woorden hoe winkels uit de treatmentgroep reageren op goede of slechte prestaties van hun tegenstanders.

We vinden dat winkels die geen kans meer maken om een bonus te winnen, niet reageren op de tussenstanden. Echter, in de groep winkels die nog kans maken op de bonus vinden we dat 1 procentpunt stijging in de relatieve prestaties van een winkel leidt tot een omzetstijging van 0,72% in de daaropvolgende week. Omdat door toeval slechts weinig winkels na de eerste week nog een kans hadden om de bonus te winnen, heeft het semi-toernooi gemiddeld genomen niet tot extra omzet geleid.

Hoofdstuk 5: prijsstructuur in eliminatietoernooien

Het veldexperiment dat in hoofdstuk 5 beschreven wordt, concentreert zich op eliminatietoernooien. Meer specifiek, we kijken naar verkoopwedstrijden over twee ronden, waarbij de slechtst presterende winkels van iedere groep na de eerste ronde afvallen, terwijl de besten een prijs winnen en doorgaan naar de tweede ronde. In de tweede ronde kunnen de overgebleven winkels opnieuw een prijs winnen. We bestuderen hoe de verdeling van het prijzengeld over de twee ronden van een dergelijk eliminatietoernooi de prestaties gedurende het toernooi beïnvloedt. De theorie voorspelt dat naarmate het prijzengeld geconcentreerd is in de tweede ronde, de prestaties in de tweede ronde verbeteren ten koste van de prestaties in de eerste ronde. Daarnaast onderzoeken we of het stimulerende effect van het toernooi afhangt van ruis in de prestatiemaatstaf. Economische theorie voorspelt dat toevallige fluctuaties in de prestaties een negatief effect hebben op de prikkel om te presteren. De reden is dat wanneer de rol van het toeval groot is, de kans dat extra inspanningen daadwerkelijk resulteren in het winnen van een prijs zeer gering is. Onze studie is de eerste die deze twee vragen onderzoekt middels een veldexperiment in een bestaande organisatie. Eerdere studies maken gebruik van bestaande data (Audas et al 2004, DeVaro 2006) of een laboratoriumexperiment (Altmann et al. 2008 en Bull 1987). Het voordeel van een veldexperiment is dat het toernooi plaatsvindt in de dagelijkse omgeving van de werknemers, terwijl door de aanwezigheid van een controlegroep toch een betrouwbare schatting van de effecten mogelijk is.

Het experiment werd uitgevoerd in dezelfde winkelketen als in hoofdstuk 4. Voor een aantal willekeurig gekozen winkels organiseren we eliminatietoernooïen over twee ronden, de overige winkels vormen de controlegroep. In beide ronden worden de deelnemende winkels ingedeeld in poules van vier vergelijkbaar presterende winkels. De helft van de winkels valt uit na de eerste ronde, de andere helft wint een prijs en gaat door naar de tweede ronde. In de tweede ronde winnen de twee best presterende winkels van iedere poule opnieuw een prijs. Als prestatiemaatstaf gebruiken we het aantal verkochte artikelen per klant.

Om het effect van de verdeling van het prijzengeld over de twee ronden vast te stellen, creëren we twee verschillende treatmentgroepen die alleen verschillen in de verdeling van het prijzengeld. In de ene treatmentgroep is de prijs identiek over de twee ronden, in de andere treatmentgroep is de prijs in de tweede ronde vier maal zo hoog als in de eerste ronde. We vinden dat de prestaties in de tweede ronde 1 procentpunt beter zijn in de groep met de hoge tweede-ronde prijs, terwijl de prestaties in de eerste ronde 0.8 procentpunt lager zijn. Deze verschillen komen weliswaar overeen met de theoretische voorspellingen, maar zijn niet statistisch significant.

Om de invloed van ruis in de prestatiemaatstaf op het stimulerende effect van het toernooi te schatten, verdelen we de treatmentgroep in twee groepen: een groep met grote fluctuaties in de historische prestaties en een groep met geringe fluctuaties. Beide groepen blijven gedurende het toernooi gescheiden, zodat alle winkels in een bepaalde poule te kampen hebben met een vergelijkbare hoeveelheid ruis. We concluderen dat variabiliteit in de historische prestaties inderdaad het stimulerende effect van het toernooi negatief beïnvloedt. Dit effect is substantieel: 30% van de winkels reageren niet op het toernooi vanwege de instabiliteit van de prestatiemaatstaf. Omdat de meerderheid van de winkels dus wel op de verkoopwedstrijd reageerden, is het totaaleffect van de actie een stijging van het aantal verkochte artikelen per klant van 1,5%.

Bibliography

Agrawal, Deepak, and Rajiv Lal (1995), Contractual Arrangements in Franchising: An Empirical Investigation, *Journal of Marketing Research*, 32(2): 213-221.

Akerlof, George A. (1982), Labor Contracts as Partial Gift Exchange, *Quarterly Journal of Economics*, 97(4): 543-569.

Altmann, Steffen, Armin Falk, and Mathias Wibral (2008), Promotions and Incentives: The Case of Multi-Stage Elimination Tournaments, *Journal of Labor Economics,* forthcoming.

Aoyagi, Masaki (2010), Information Feedback in a Dynamic Tournament, *Games and Economic Behavior*, forthcoming.

Arbak, Emrah, and Laurence Kranich (2005), Can Wages Signal Kindness?, GATE Working Paper 05-11.

Arellano, M., and S. Bond (1991), Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations, *Review of Economic Studies,* vol. 58(2), pp. 277-297.

Asch, Beth J. (1990), Do Incentives Matter?, The Case of Navy Recruiters, *Industrial and Labor Relations Review*, 43(3): 89-106.

Audas, Rick, Tim Barmby, and John Treble (2004), Luck, Effort, and Reward in an Organizational Hierarchy, *Journal of Labor Economics*, 22(2): 379-396.

Auriol, Emmanuelle, and Régis Renault (2008), Status and incentives, *RAND Journal of Economics*, 39(1): 305-326.

Azmat, Ghazala, and Nagore Iriberri (2010), The Importance of Relative Performance Feedback Information: Evidence from a Natural Experiment using High School Students, *Journal of Public Economics*, 94(7-8): 435-452.

Babin, Barry, and James S. Boles (1996), The Effects of Perceived Co-Worker Involvement and Supervisor Support on Service Provider Role Stress, Performance and Job Satisfaction, *Journal of Retailing*, 72(1): 57-75.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul (2009), Team Incentives: Evidence from a Field Experiment, Mimeo, University College London.

Barro, Robert (1974), Are Government Bonds Net Wealth?, *Journal of Political Economy*, 82(6): 1095-1117.

Bartling, Björn, Ernst Fehr, and Klaus M Schmidt (2011), Screening, Competition, and Job Design: Economic Origins of Good Jobs" *American Economic Review*, forthcoming.

Battigalli, Pierpaolo, and Martin Dufwenberg (2009), Dynamic Psychological Games, *Journal of Economic Theory*, 144(1): 1-35.

Becker, Brian E., and Mark A. Huselid (1992), The Incentive Effects of Tournament Compensation Systems, *Administrative Science Quarterly*, 37(2): 336-350.

Becker, Gary (1974), A Theory of Social Interactions, *Journal of Political Economy*, 82(6): 1063-1094.

Belfield, Richard, and David Marsden (2003), Performance Pay, Monitoring Environments, and Establishment Performance'. *International Journal of Manpower*, 24(4): 452-471.

Bellemare, Charles, and Bruce Shearer (2011), On the Relevance and Composition of Gifts within the Firm: Evidence from Field Experiments, *International Economic Review*, forthcoming

Bernheim, B.D., and Oded Stark (1988): Altruism within the Family Reconsidered: Do Nice Guys Finish Last?, *American Economic Review*, 78(5), 1034-1045.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004), How Much Should We Trust Differences-in-Differences Estimates?, *Quarterly Journal of Economics*, 119(1): 249-275.

Besley, Timothy J., and Maitreesh Ghatak (2005), Competition and Incentives with Motivated Agents, *American Economic Review*, 95(3): 616-636.

Besley, Timothy, and Maitreesh Ghatak (2008), Status Incentives, *American Economic Review*, 98(2): 206–211.

Bewley, Truman F. (1999), Why Wages Don't Fall during a Recession, Harvard University Press.

Blanes i Vidal, Jordi, and Mareike Nossol (2009), Tournaments without Prizes: Evidence from Personnel Records, *Management Science*, forthcoming.

Blau, Peter M. (1964), Exchange and Power in Social Life, John Wiley, New York.

Bognanno, Michael L. (2001), Corporate Tournaments, Journal of Labor Economics, 19(2): 290-315.

Borzaga, Carlo, and Sara Depedri (2005), Interpersonal Relations and Job Satisfaction: Some Empirical Results, In: Benedetto Gui and Robert Sugden (Eds.), Economics and Social Interaction; Accounting for Interpersonal Relations, Cambridge: Cambridge University Press.

Bouwens, Jan, and Peter Kroos (2011), Target Ratcheting and Effort Reduction, *Journal of Accounting and Economics*, 51(1-2), 171-185.

Brandts, Jordi, and Gary Charness (2003), Truth or Consequences: An Experiment, *Management Science*, 49(1): 116-130.

Bull, Clive, Andrew Schotter, and Keith Weigelt (1987), Tournaments and Piece Rates: An Experimental Study, *Journal of Political Economy*, 95(1): 1-33.

Cadsby, C. Bram, Fei Song, and Francis Tapon (2010), Are You Paying Your Employees to Cheat? An Experimental Investigation, *B.E. Journal of Economic Analysis & Policy*, 10(1), article 35.

Camerer, Colin, and Richard H. Thaler (1995), Anomalies: Ultimatums, Dictators, and Manners, *The Journal of Economic Perspectives*, 9(2): 209-219.

Campbell, Carl M., and Kunal S. Kamlani (1997), The Reasons for Wage Rigidity: Evidence from a Survey of Firms, *Quarterly Journal of Economics*, 112(3): 759-789.

Carmichael, H. Lorne (1983), The Agents-Agents Problem: Payment by Relative Output, *Journal of Labor Economics*, 1(1): 50-60.

Casas-Arce, Pablo, and F. Asis Martinez-Jerez (2009), Relative Performance Compensation, Contests, and Dynamic Incentives, *Management Science*, 55(8): 1306-1320.

Charness, Gary, and Peter J. Kuhn (2010), Lab Labor, What Can Labor Economist Learn from the Lab?, Forthcoming in: Orley Ashenfelter and David Card (Eds.), Handbook of Labor Economics, Vol. 4, Amsterdam: Elsevier.

Charness, Gary, Peter Kuhn, and Marie-Claire Villeval (2010), Competition and the Ratchet Effect, mimeo.

Charness, Gary, and David Levine (2007), Intention and Stochastic Outcomes: An Experimental Study, *Economic Journal*, 117(522): 1051-1072.

Cho, In-Koo, and David M. Kreps (1987), Signaling Games and Stable Equilibria, *Quarterly Journal of Economics*, 102(2): 179-221.

Cooper, David, John Kagel, Wei Lo, and Qing Liang Gu (1999), Gaming Against Managers in Incentive Systems: Experimental Results with Chinese Students and Chinese Managers, *American Economic Review*, 89(4): 781-804.

Cropanzano, Russel, and Marie S. Mitchell (2005), Social Exchange Theory: An Interdisciplinary Review, *Journal of Management*, 31(6): 874-900.

Delfgaauw, Josse and Robert Dur (2007), Signaling and Screening of Workers' Motivation, *Journal of Economic Behavior & Organization*, 62(4): 605-624.

Delfgaauw, Josse, Robert Dur, Arjan Non, and Willem Verbeke (2010), Dynamic Incentive Effects of Relative Performance Pay: A Field Experiment, Tinbergen Institute Discussion Paper 10-124/1.

Delfgaauw, Josse, Robert Dur, Joeri Sol, and Willem Verbeke (2009), Tournament Incentives in the Field: Gender Differences in the Workplace, Tinbergen Institute Discussion Paper 09-061/1.

Demski, Joel S., and David E.M. Sappington (1991), Resolving Double Moral hazard Problems with Buyout Agreements, *RAND Journal of Economics*, 22(2): 232-240.

DeVaro, Jed (2006), Internal Promotion Competition in Firms, *RAND Journal of Economics*, 37(3): 521-542.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2009), Homo Reciprocans: Survey Evidence on Behavioural Outcomes, *Economic Journal,* 119: 592-612.

Dufwenberg, Martin, and Georg Kirchsteiger (2004), A Theory of Sequential Reciprocity, *Games and Economic Behavior*, 47(2): 268-298.

Dur, Robert (2009), Gift Exchange in the Workplace: Money or Attention?, *Journal of the European Economic Association*, 07(2-3): 550-560.

Dur, Robert, Arjan Non, and Hein Roelfsema (2008), Reciprocity and Incentive Pay in the Workplace, Tinbergen Discussion Paper No 08-080/1

Dur, Robert, Arjan Non, and Hein Roelfsema (2010), Reciprocity and Incentive Pay in the Workplace, *Journal of Economic Psychology*, 31(4): 676-686.

Dur, Robert, and Hein Roelfsema (2010), Social Exchange and Common Agency in Organizations, *Journal of Socio-Economics*, 39(1): 55-63.

Dur, Robert, and Joeri Sol (2010): "Social Interaction, Co-Workers Altruism, and Incentives, *Games and Economic Behavior*, 69(2): 293-301.

Ederer, Florian (2010), Feedback and Motivation in Dynamic Tournaments, *Journal of Economics & Management Strategy*, 19(3): 733-769.

Ehrenberg, Ronald G., and Michael L. Bognanno (1990), Do Tournaments Have Incentive Effects?, *Journal of Political Economy*, 98(6): 1307-1324.

Eisenberger, R., Huntington, R., Hutchison, S., & Sowa, D. (1986). "Perceived organizational support, *Journal of Applied Psychology*, 71(3), 500-507.

Englmaier, Florian and Stephen Leider (2008), Contractual and Organizational Structure with Reciprocal Agents, CESifo Working Paper 2415.

Erev, I., G. Bornstein, and G. Rachely (1993), Constructive Intergroup Competition as a Solution to the Free Rider Problem: A Field Experiment, *Journal of Experimental Social Psychology*, 29(6): 463-478.

Eriksson, Tor (1999), Executive Compensation and Tournament Theory, *Journal of Labor Economics*, 17(2): 262-280.

Falk, Armin, Ernst Fehr, and Urs Fischbacher (2003), On the Nature of Fair Behavior, *Economic Inquiry*, 41(1): 20-26.

Falk, Armin, Ernst Fehr, and Urs Fischbacher (2008), Testing theories of fairness – Intentions matter, *Games and Economic Behavior*, 62(1): 287-303.

Falk, Armin, and Urs Fischbacher (2006), A Theory of Reciprocity, *Games and Economic Behavior*, 54(2): 293-315.

Falk, Armin, and Michael Kosfeld (2006), The Hidden Costs of Control, *American Economic Review*, 96(5), 1611–30.

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.

Fehr, Ernst, and Simon Gächter (1998), Reciprocity and economics: the economic implications of Homo Reciprocans". *European Economic Review*, 42(3): 845-859.

Fehr, Ernst, and Simon Gächter (2000), Fairness and Retaliation: The Economics of Reciprocity, *Journal of Economic Perspectives,* 14(3): 159-181.

Fehr, Ernst, and Simon Gächter (2002), Do incentive contracts undermine voluntary cooperation?, Working Paper No. 34, Institute for Empirical Research in Economics, University of Zürich.

Fehr, Ernst, George Kirchsteiger and Arno Riedl (1993), Does Fairness Prevent Market Clearing?, An Experimental Investigation, *Quarterly Journal of Economics,* 108(2): 437-459.

Fershtman, Chaim, and Uri Gneezy (2011), The Trade-off between Performance and Quitting in High-Power Tournaments, *Journal of the European Economic Association*, 9(2): 318-336.

Fischbacher, Urs, and Simon Gächter (2010) "Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments, *American Economic Review*, 100(1): 541–56.

Fischbacher, Urs, Simon Gächter, and Ernst Fehr (2001) "Are People Conditionally Cooperative?, Evidence from a Public Goods Experiment, *Economics Letters*, 71(3): 397–404.

Frank, Douglas, and Tomasz Obloj (2011), Reference Points and Organizational Performance: Evidence from Retail Banking, mimeo, INSEAD.

Freeman, Richard B, and Alexander M. Gelber (2010), Prize Structure and Information in Tournaments: Experimental Evidence, *American Economic Journal: Applied Economics*, 2(1): 149-164.

Freeman, Richard B., Douglas Kruse and Joseph Blasi (2008), The Same Yet Different: Worker reports on Labour Practices and Outcomes in a Single Firm across Countries, *Labour Economics,* 15(4): 749-770.

Frey, Bruno S., and Susanne Neckermann (2008), Awards - A View from Psychological Economics, *Journal of Psychology*, 216(4): 198-208.

Fu, Qiang, and Jingfeng Lu (2009), The Optimal Multi-stage Contest, *Economic Theory,* forthcoming.

Gächter, Simon and Christian Thöni (2010), Social Comparison and Performance: Experimental Evidence on the Fair Wage-Effort Hypothesis, IZA Working paper No. 4687.

Gershkov, Alex, and Motty Perry (2009), Tournaments with Midterm Reviews, *Games and Economic Behavior*, 66: 162-190.

Gibbs, Michael (1995), Incentive Compensation in a Corporate Hierarchy, *Journal of Accounting and Economics,* 19(2-3): 247-277.

Gibbs, Michael (1996), Promotions and Incentives, mimeo, University of Chicago.

Gittell, Jody Hoffer (2003), The Southwest Airlines Way; Using the Power of Relationships to Achieve High Performance, New York: McGraw-Hill.

Gneezy, Uri, and John A. List (2006), Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments, *Econometrica,* 74(5): 1365-1384.

Goltsman, Maria and Arijit Mukherjee (2010), Interim Performance Feedback in Multistage Tournaments: The Optimality of Partial Disclosure, Journal of Labor Economics, forthcoming.

Graafland, Johan J. and B.A. Rutten (2004), Perceived Organisational Support and Profitability, CentER Discussion Paper 2004-45.

Green, Jerry R., and Nancy L. Stokey (1983), A Comparison of Tournaments and Contracts, *Journal of Political Economy,* 91(3): 349-364.

Groh, Christian, Benny Moldovanu, Aner Sela, and Uwe Sunde, (2010), Optimal Seedings in Elimination Tournaments, *Economic Theory,* forthcoming.

Grund, Christian, and Dirk Sliwka (2005), Envy and compassion in tournaments, *Journal of Economics and Management Strategy,* 14(1): 187-207.

Grund, Christian, and Dirk Sliwka (2010), Evidence on performance pay and risk aversion, *Economics Letters*, 106(1): 8-11.

Gul, Faruk, and Wolfgang Pesendorfer (2010), Interdependent Preference Models as a Theory of Intentions, Princeton University Working Paper.

Gupta, Srabana and Richard E. Romano (1998), Monitoring the Principal with Multiple Agents,RAND *Journal of Economics*, 29(2): 427- 442.

Harbring, Christine, and Bernd Irlenbusch (2003), An Experimental Study on Tournament Design, *Labour Economics*, 10(4): 443-464.

Harrison, Glenn, and John A. List (2004), Field Experiments, *Journal of Economic Literature*, 42(4): 1009-1055.

Hennig-Schmidt, Heike, Bettina Rockenbach, and Abdolkarim Sadrieh (2010), In Search of Workers' Real Effort Reciprocity – A Field and a Laboratory Experiment, *Journal of the European Economic Association*, 8(4): 817-837.

Höchtl, Wolfgang, Rudolf Kerschbamer, Rudi Stracke, and Uwe Sunde (2011), Effort versus Sorting: That's the Question. Two-Stage Elimination Tournament in Theory and Experiment, mimeo, University of Innsbruck.

Homans, G.C. (1958), Social behavior as exchange". *American Journal of Sociology*, 63(6): 597-606.

Huck, Steffen, and Jean-Robert Tyran (2007). "Reciprocity, Social Ties, and Competition in Markets for Experience Goods, *Journal of Socio-Economics*, 36(2): 191-203.

Kamdar, Dishan and Linn Van Dyne (2007), The Joint Effects of Personality and Workplace Social Exchange Relationships in Predicting Task Performance and Citizenship Performance, *Journal of Applied Psychology*, 92(5): 1286-1298.

Knoeber, Charles R., and Walter N. Thurman (1994), Testing the Theory of Tournaments: An Empirical Analysis of Broiler Production, *Journal of Labor Economics*, 12(2): 155-179.

Kosfeld, Michael, and Susanne Neckermann (2010), Getting More Work for Nothing?, Symbolic Awards and Worker Performance, *American Economic Journal: Microeconomics*, forthcoming.

Kosfeld, Michael, and Ferdinand A. von Siemens (2009), Worker Self-Selection and the Profits from Cooperation, *Journal of the European Economic Association*, 7(2-3): 573-582.

Kosfeld, Michael, and Ferdinand von Siemens (2011), Competition, Cooperation, and Corporate Culture, *The RAND Journal of Economics*, 42(1): 23-43.

Kube, Sebastian, Michel André Maréchal, and Clemens Puppe (2008), The Currency of Reciprocity – Gift-Exchange in the Workplace, Working paper, Institute for Empirical Research in Economics, University of Zurich.

Lafontaine, Francine (1992), Agency Theory and Franchising: Some Empirical Results, *The RAND Journal of Economics,* 23(2): 263-283.

Lazear, Edward P. (1989), Pay Equality and Industrial Politics, *Journal of Political Economy*, 97(3): 561-580.

Lazear, Edward P. (2000), Performance Pay And Productivity, *American Economic Review,* 90(5): 1346-1361.

Lazear, Edward P., and Paul Oyer (2009), Personnel Economics. Forthcoming in: Robert Gibbons and John Roberts (Eds.), Handbook of Organizational Economics. Princeton: Princeton University Press.

Lazear, Edward P., and Sherwin Rosen (1981), Rank-Order Tournaments as Optimum Labor Contracts, *Journal of Political Economy*, 89(5): 841-864.

Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent (2009), Performance Pay and Wage Inequality, *The Quarterly Journal of Economics,* 124 (1): 1-49.

Levine, David K. (1998), Modeling Altruism and Spitefulness in Experiments, *Review of Economic Dynamics*, 1(3): 593-622.

Ley, Eduardo (1997), Optimal Provision of Public Goods with Altruistic Individuals, Economic Letters, 54(1): 23-27.

Frank, Douglas, and Tomasz Obloj (2011), Reference Points and Organizational Performance: Evidence from Retail Banking, mimeo, INSEAD.

Lim, Noah, Michael J. Ahearne, and Sung H. Ham (2009), Designing Sales Contests: Does the Prize Structure Matter?, *Journal of Marketing Research*, vol. 46(3): 356-371.

McCue, Kristin (1996), Promotions and Wage Growth, *Journal of Labor Economics,* 14(2): 175-209.

Minkler, Lanse (2004), Shirking and motivation in firms: survey evidence on worker attitudes, *International Journal of Industrial Organisation*, 22(6): 863-884.

Moldovanu, Benny, Aner Sela, and Xianwen Shi (2007), Contests for Status, *Journal of Political Economy,* 115(2): 338–363.

Murphy, Kevin J. (1985), Corporate Performance and Managerial Remuneration: An Empirical Analysis, *Journal of Accounting and Economics,* 7(1-3): 11-42.

Nadler, David A., Michael L. Tushman (1997), Competing by Design: The Power of Organizational Architecture, New York, NY: Oxford University Press.

Nagin, Daniel S., James B. Rebitzer, Seth Sanders, and Lowell J. Taylor (2002), Monitoring, Motivation, and Management: The Determinants of Opportunistic Behavior in a Field Experiment, *American Economic Review,* 92(4): 850-873.

Nalebuff, Barry J., and Joseph E. Stiglitz (1983), Prizes and Incentives: Towards a General Theory of Compensation and Competition, Bell *Journal of Economics*, 14(1): 21-43.

Nöldeke, Georg, and Klaus Schmidt (1995), Option Contracts and Renegotiation: A Solution to the Hold-Up Problem, *The RAND Journal of Economics*, 26(2): 163-179.

Oyer, Paul (1998), Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality, *Quarterly Journal of Economics*, 113(1): 149-185.

Pazy, Asya, and Yoav Ganzach (2009), Pay Contingency and the Effects of Perceived Organizational and Supervisor Support on Performance and Commitment, *Journal of Management*, 35(4): 1007-1025.

Pfeffer, Jeffrey (1998), The Human Equation: Building Profits by Putting People First, Boston, MA: Harvard Business School Press.

Piekkola, Hannu (2005), Performance-related pay and firm performance in Finland, *International Journal of Manpower,* 26(7/8): 619-635.

Prendergast, Canice (1999), The Provision of Incentives in Firms, *Journal of Economic Literature,* 37(1): 7-63.

Rabin, Matthew (1993), Incorporating Fairness into Game Theory and Economics, *American Economic Review,* 83(5): 1281-1302.

Rhoades, Linda, and Robert Eisenberger (2002), Perceived Organizational Support: A Review, *Journal of Applied Psychology*, 87(4): 698-714.

Roethlisberger, Fritz J. and William Dickson (1939), Management and the Worker, Cambridge, MA: Harvard University Press.

Rosen, Sherwin (1986), Prizes and Incentives in Elimination Tournaments, *American Economic Review,* 76(4): 701-715.

Royal Dutch Shell (2009), Annual Report and Form 20-F for the Year ended December 31, 2009.

Shanock, Linda R., and Robert Eisenberger (2006), When Supervisors Feel Supported: Relationships With Subordinates' Perceived Supervisor Support, Perceived Organizational Support, and Performance, *Journal of Applied Psychology*, 91(3): 689–695.

Shearer, Bruce (2004), Piece Rates, Fixed Rates, and Incentives: Evidence from a Field Experiment, *Review of Economic Studies,* 71(2): 513-534.

Sheremeta, Roman M. (2010), Experimental Comparison of Multi-stage and One-stage Contests, *Games and Economic Behavior,* 68(2): 721-747.

Sliwka, Dirk (2007), Trust as a Signal of a Social Norm and the Hidden Costs of Incentive Schemes, *American Economic Review,* 97(3): 999-1012.

Soberman, David A., (2003). "Simultaneous Signaling and Screening with Warranties, *Journal of Marketing Research,* 40 (2): 176-209.

Sol, Joeri, (2010), Peer Evaluation: Incentives and Co-Worker Relations, Tinbergen Institute Discussion Paper 10-055/1.

Vogt, H. (1983), Unimodality of Differences, *Metrika*, vol. 30(1), pp. 165-170.

Waldman, Michael (2008), Theory and Evidence in Internal Labor Markets. Forthcoming in: Robert Gibbons and John Roberts (Eds.), Handbook of Organizational Economics. Princeton: Princeton University Press.

Wayne, Sandy J., Lynn M. Shore, and Robert C. Liden (1997), Perceived Organizational Support and Leader-Member Exchange: A Social Exchange Perspective, *Academy of Management Journal*, 40(1): 82-111.

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam and Vrije Universiteit 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:

460 O.E. JONKEREN, Adaptation to Climate Change in Inland Waterway Transport.
461 S.C. GO, Marine Insurance in the Netherlands 1600-1870, A Comparative Institutional Approach.
462 J. NIEMCZYK, Consequences and Detection of Invalid Exogeneity Conditions.
463 I. BOS, Incomplete Cartels and Antitrust Policy: Incidence and Detection
464 M. KRAWCZYK, Affect and risk in social interactions and individual decision-making.
465 T.C. LIN, Three Essays on Empirical Asset Pricing.
466 J.A. BOLHAAR, Health Insurance: Selection, Incentives and Search.
467 T. FARENHORST-YUAN, Efficient Simulation Algorithms for Optimization of Discrete Event Based on Measure Valued Differentiation.
468 M.I. OCHEA, Essays on Nonlinear Evolutionary Game Dynamics.
469 J.L.W. KIPPERSLUIS, Understanding Socioeconomic Differences in Health: An Economic Approach.
470 A. AL-IBRAHIM, Dynamic Delay Management at Railways: A Semi-Markovian Decision Approach.
471 R.P. FABER, Prices and Price Setting.
472 J. HUANG, Education and Social Capital: Empirical Evidences from Microeconomic Analyses.
473 J.W. VAN DER STRAATEN, Essays on Urban Amenities and Location Choice.
474 K.M. LEE, Filtering Non Linear State Space Models: Methods and Economic Applications.
475 M.J. REINDERS, Managing Consumer Resistance to Innovations.
476 A. PARAKHONYAK, Essays on Consumer Search, Dynamic Competition and Regulation.
477 S. GUPTA, The Study of Impact of Early Life Conditions on Later Life Events: A Look Across the Individual's Life Course.
478 J. LIU, Breaking the Ice between Government and Business: From IT Enabled Control Procedure Redesign to Trusted Relationship Building.
479 D. RUSINOVA, Economic Development and Growth in Transition Countries.
480 H. WU, Essays on Top Management and corporate behavior.
481 X. LIU, Three Essays on Real Estate Finance
482 E.L.W. JONGEN, Modelling the Impact of Labour Market Policies in the Netherlands
483 M.J. SMIT, Agglomeration and Innovations: Evidence from Dutch Microdata

484 S.VAN BEKKUM, What is Wrong With Pricing Errors? Essays on Value Price Divergence
485 X. HU, Essays on Auctions
486 A.A. DUBOVIK, Economic Dances for Two (and Three)
487 A.M. LIZYAYEV, Stochastic Dominance in Portfolio Analysis and Asset Pricing
488 B. SCHWAAB, Credit Risk and State Space Methods
489 N. BASTÜRK, Essays on parameter heterogeneity and model uncertainty
490 E.GUTIÉRREZ PUIGARNAU, Labour markets, commuting and company cars
491 M.W. VORAGE, The Politics of Entry
492 A.N. HALSEMA, Essays on Resource Management: Ownership, Market Structures and Exhaustibility
493 R.E. VLAHU, Three Essays on Banking
494 N.E. VIKANDER, Essays on Teams and the Social Side of Consumption
495 E. DEMIREL, Economic Models for Inland Navigation in the Context of Climate Change
496 V.A.C. VAN DEN BERG, Congestion pricing with Heterogeneous travellers
497 E.R. DE WIT, Liquidity and Price Discovery in Real Estate Assets
498 C. LEE, Psychological Aspects of the Disposition Effect: An Experimental Investigation
499 MHA. RIDHWAN, Regional Dimensions of Monetary Policy in Indonesia
500 J. GARCÍA, The moral herd: Groups and the Evolution of Altruism and Cooperation
501 F.H. LAMP, Essays in Corporate Finance and Accounting
502 J. SOL, Incentives and Social Relations in the Workplace
503 A.I.W. HINDRAYANTO, Periodic Seasonal Time Series Models with applications to U.S. macroeconomic data
504 J.J. DE HOOP, Keeping Kids in School: Cash Transfers and Selective Education in Malawi
505 O. SOKOLINSKIY, Essays on Financial Risk: Forecasts and Investor Perceptions
506 T. KISELEVA, Structural Analysis of Complex Ecological Economic Optimal Management Problems
507 U. KILINC, Essays on Firm Dynamics, Competition and Productivity
508 M.J.L. DE HEIDE, R&D, Innovation and the Policy Mix
509 F. DE VOR, The Impact and Performance of Industrial Sites: Evidence from the Netherlands

Economists have traditionally stressed the importance of monetary incentives for employee motivation, to the exclusion of other relevant motives. One of these motives is reciprocity, meaning that people respond in kind to hostile and friendly acts. Reciprocity implies that workers are more willing to go the extra mile when they perceive their employer as generous. The first part of this thesis explores theoretically how monetary incentives influence, and are influenced by, worker reciprocity. For example, assuming workers are reciprocal towards management attention, I compare promotion incentives with individual performance pay. The second part of this thesis presents the results of two field experiments conducted in a Dutch retail chain. The experiments consist of sales contests that are explicitly designed to test predictions from tournament theory regarding feedback, prize structure and fluctuations in the performance measure.

Arjan Non (1983) graduated in Economics at the Erasmus University Rotterdam in 2006. Afterwards, he joined the Tinbergen Institute as a PhD-student at the Department of Economics at Erasmus University. Currently, he just started as a post-doctoral researcher at Maastricht University & The Research Centre for Education and the Labour Market (ROA). A part of this thesis has been published in the Journal of Economic Psychology.

Thela Thesis
Academic Publishing Services