Social Preferences and Incentives in the Workplace

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
Steffen Altmann
aus Villingen-Schwenningen

Bonn 2009

Dekan:	Prof. Dr. Christian Hillgruber
Erstreferent:	Prof. Dr. Armin Falk
Zweitreferent:	Prof. Paul Heidhues, Ph.D.

Tag der mündlichen Prüfung: 28.07.2009

Diese Dissertation ist auf dem Hochschulschriftenserver der ULB Bonn (http://hss.ulb.uni-bonn.de/diss_online) elektronisch publiziert.

Acknowledgments

First and foremost, I am grateful to Armin Falk for his excellent advice, inexhaustible support and firm guidance throughout the last years. When writing my diploma thesis, I came across his dissertation which awaked my interest in the exciting field of "behavioral" labor economics. I was impressed about the originality and creativity with which he addressed important questions in an innovative way. Working with him, however, exceeded all my expectations.

Paul Heidhues, Matthias Kräkel and several other faculty members of the Bonn Graduate School of Economics (BGSE) have always been ready to discuss ideas and provided very helpful comments during various stages of this dissertation.

I want to thank my co-authors Johannes Abeler, Thomas Dohmen, David Huffman, Sebastian Kube, and Matthias Wibral. Collaborating, discussing, conducting experiments, analyzing data, and writing papers with them has been great fun.

Financial support from the DFG (through GRK 629), the BGSE, the Institute for Empirical Research in Economics at the University of Zurich, and the Institute for the Study of Labor (IZA) is gratefully acknowledged. I especially thank Urs Schweizer, Georg Nöldeke and Jürgen von Hagen for their continuous engagement in making the BGSE such a stimulating place to start doing economic research.

Even the BGSE, however, could not be such a great place without great fellow students. Johannes Abeler, An Chen, Oliver Fries, Jürgen Gaul, Sotiris Georganas, Simon Junker, Lars Koch, Julia Nafziger, Michal Paluch, Marc Schiffbauer, David Schröder, Marcus Sonntag, Xia Su, Matthias Wibral, and Michael Zähringer were the best 'year 2003 graduate class' I could imagine.

IZA provided and continues to provide a very inspiring atmosphere for doing research. I am glad to be part of such a great team. In particular, I would like to thank Franziska Tausch for excellent research assistance on various projects of this dissertation. I owe great thanks to BonnEconLab, and especially to Heike Hennig-Schmidt whose effort and support allows me to conduct economic experiments in a "perfectly controlled" environment. I also want to thank the 792 anonymous (!) subjects whose behavior will be discussed in the following pages.

Many people inside and outside the University of Bonn made sure that life "after research" has also been a great pleasure. Apart from my fellow students, these people include Almira Buzaushina, Zeno Enders, Sebastian Goerg, Gesine Güllner, Johannes Kaiser, Carolin Kieß, Simon Vehovec, Gari Walkowitz, and Markus Zink.

Most of all, I thank my family for their encouragement, patience, and unconditional support in everything I do.

Contents

List of Figures

List of Tables

Introduction

Most economic exchanges are not based on fully contingent, explicit contracts, but rather rely on informal agreements that specify the contracting parties' obligations only imprecisely. There are several reasons why this is the case. First, the costs of writing fully contingent contracts can be too high. Moreover, the contracting parties often cannot or do not foresee all future contingencies at the time when they enter the trade agreement. Finally, many economic interactions are characterized by private information or opportunities to take hidden actions. It can thus be impossible—or at least prohibitively costly—to monitor and verify the fulfillment of a trading partner's contractual obligations.

In the labor market, contractual incompleteness is particularly widespread. Employment contracts in many occupations are rather general agreements that specify far from all aspects of an employment relation. Contracts often only stipulate a fixed salary and a required working time but do not, e.g., specify precisely which tasks an employee has to fulfill. Moreover, pay is rarely tied explicitly to performance. One explanation for this lies in the fact that work effort and performance are inherently difficult to verify. In addition, research in contract and game theory has identified several circumstances where it might not be necessary or even undesirable for an employer to connect pay more closely to performance although it would, in principle, be possible.

First, employment relations are typically "repeated interactions" over a long time horizon. Therefore, there is scope for self-enforcing relational contracts that are sustained by the value of future interaction (MacLeod and Malcomson 1989, Baker et al. 1992). Second, and based on a related argument, employees have concerns for their future career. If employees are looking for a better paid job in the future, either within the firm (Lazear and Rosen 1981, Rosen 1986) or in other firms (Fama 1980, Holmström 1999, Gibbons and Murphy 1992), they have strong implicit incentives to perform well, although pay is not tied to current performance. Finally, an employer might not want to use high-powered pay-for-performance incentives if he aims at inducing the employee to engage in multiple tasks, but can only verify performance in a subset of these tasks (Holmström and Milgrom 1991).

However, there could also be psychological reasons why employers do not want or need to use high-powered incentive schemes. In recent years, a growing body of literature in psychology and behavioral economics has enhanced our understanding of the psychological foundations of incentives and has demonstrated that employee behavior does not only depend on the relationship between measured performance and pay. Rather, pay-for-performance schemes can have unintended, dysfunctional consequences if workers are concerned about horizontal or vertical equity (Adams 1965), are intrinsically motivated (Deci 1971, Gneezy and Rustichini 2000), or dislike being controlled (Falk and Kosfeld 2006). Moreover, it has been shown that social preferences can be able to mitigate moral-hazard problems. Especially reciprocity, i.e., the willingness to reward kind actions and punish unkind ones even at a cost, has proven to be highly effective in eliciting work effort under contractual incompleteness (Fehr et al. 1997, Fehr and Falk 2002).

A deeper understanding of the interaction between monetary incentives and non-monetary motivations has allowed to advance economic models of labor market behavior, e.g., by analyzing the consequences of social preferences (Fehr and Schmidt 1999), self-control problems (O'Donoghue and Rabin 1999), or reference points such as relative income concerns or personal income targets (Neumark and Postlewaite 1998, Köszegi and Rabin 2006). This has also shed new light on the effectiveness of important labor market institutions such as minimum wage laws (Falk et al. 2006) or employment protection legislation (Falk, Huffman and MacLeod 2008).

This dissertation aims at further enhancing our understanding of behavior under contractual incompleteness, with a particular focus on employment relationsships. All chapters analyze situations where explicit pay-for-performance contracts are not feasible, and therefore gains from trade can only be realized if trading parties successfully make use of implicit monetary or non-monetary incentives. All of the four chapters report evidence from laboratory experiments. Controlled laboratory and field experiments have rapidly emerged as a vital component of research in labor economics. For the questions addressed in this dissertation, the use of experimental methods allows identifying direct causal effects since treatment conditions (e.g., incentive schemes) can be varied exogenously.

Chapters 1–3 concentrate on interactions that are characterized by moral hazard. It is analyzed whether and how trust and social preferences can help mitigating these problems. Chapter 1 studies the intrapersonal relationship between trust and reciprocity and explores to what extent heterogeneity in individuals' reciprocal inclination can account for commonly observed differences in trusting behavior. Chapter 2 deals with the question how horizontal fairness concerns impact the effectiveness of incentive schemes that are based on reciprocal gift exchange. Chapter 3 analyzes whether two common features of implicit contracts—gift exchange and relational contracting—can lead to involuntary unemployment. Chapter 4 studies incentives provided through promotion competitions. More precisely, it addresses the question how people behave in multi-stage elimination tournaments in comparison to simple, one-stage promotion contests.

All chapters focus on topics that are vital for understanding labor market institutions and behavior in employment relations. They are thus directly relevant for questions studied in personnel and labor economics. However, the chapters analyze the impact of monetary and non-monetary incentives on behavior, more generally. The relevance of the results presented below thus reaches beyond labor market applications. For instance, trust (studied in Chapter 1) is involved in essentially all economic transactions. Mutual trust between trading parties reduces the cost of economic exchanges, thus making it easier to realize gains from trade. Ultimately, a high level of trust might therefore be conducive to economic growth (Knack and Keefer 1997). Social comparison and equity concerns, which are driving forces for the results in Chapter 2, also influence consumption patterns and subjective well-being and therefore have potentially wide ranging welfare implications (Frank 1999, Layard 2005). The results presented in Chapter 3 potentially also provide new insights into other markets where long-term relationships and relational contracts are common features of market interaction, e.g., lender-borrower relations in the credit market (Boot 2000). Finally, competition in tournaments and other contest-like environments is highly relevant in many fields of economic and social life, such as R & D races (Taylor 1995), politics (Hillman and Riley 1989), sports (Ehrenberg and Bognanno 1990), or broiler production (Knoeber and Thurman 1994).

In the first chapter of this dissertation, we analyze the intrapersonal relationship between trust and reciprocity. While obviously trust and reciprocity are connected in the sense that the expected returns to trust are higher when interacting with a reciprocal trading partner, surprisingly little is known about whether individuals with different degrees of reciprocal inclinations also differ in their trusting behavior. In other words, do reciprocal persons trust more or less compared to more selfish ones?

To study this question, we employ a variant of the trust game (Berg et al. 1995) that allows us to measure both variables on an individual basis. Our main result shows a strong and positive relation between a person's reciprocity and her trusting behavior. Reciprocal players exhibit much higher levels of trust than more selfish individuals. This effect is highly robust, e.g., if we control for the influence of gender or risk attitudes.

Substantial amounts of trust, but also a high degree of heterogeneity in trusting behavior are key findings in empirical trust research. This variability raises the question whether a person's preferences or individual characteristics are systematically related to her behavior in situations involving trust. Recently, several determinants of trust have been identified, including risk attitudes (Dohmen et al. 2006), fear of betrayal (Bohnet and Zeckhauser 2004), gender (Buchan et al. 2008), or individuals' genetic predisposition (Reuter et al. 2008). Our findings further enhance understanding of differences in trusting behavior, showing that a person's reciprocal inclination and, hence, her own trustworthiness is strongly related to trust.

Exploring the link between an individual's trusting behavior and her reciprocal inclination is also of high importance for the evaluation and further development of theories that incorporate social preferences. Most of these theories assume—at least implicitly—a certain connection between the two variables. Some of the most prominent models (e.g., Fehr and Schmidt 1999, Bolton and Ockenfels 2000) predict that individuals who are more reciprocal or inequity averse ceteris paribus trust *less* than others in the trust game. The intuition for this result is simple: a selfish sender just suffers from the loss of her investment if the receiver sends back too little, whereas a fair-minded sender experiences additional distress because his trust has been exploited. Our finding of a strong and *positive* relationship between trust and reciprocity suggests that this intuition may be too simplistic and that models of social preferences should be extended, e.g., to account for non-common expectations of different types.

In Chapter 2, we study a situation where trust and reciprocity could be particularly conducive to increase efficiency of economic interactions—gift exchange in employer-employee relationships. In recent years, a vast body of literature has stressed the importance of reciprocity for mitigating moral hazard problems of incomplete contracts: since many agents repay a gift in the form of higher wages by providing higher efforts, effort can be elicited under incomplete contracts even in one-shot situations where no future gains can be expected (e.g., Akerlof 1982, Fehr et al. 1997, Hannan et al. 2002, Fehr and Falk 2002, Maximiano et al. 2007).

The potential of gift exchange as a contract enforcement device, however, is likely to depend on the institutions that shape the employment relation, above all the mode of payment. A key question in this context is how to treat agents relative to each other as this affects the perceived fairness of a pay scheme. In Chapter 2, we study this question by analyzing two important fairness principles: horizontal equality and equity.

The specific wage institution we consider is wage equality. Paying equal wages to workers on the same level of a hierarchy is common practice in many firms (e.g., Medoff and Abraham 1980, Baker et al. 1988). Several reasons for equal wages have been brought forward, amongst them increased peer monitoring (Knez and Simester 2001) and lower transaction costs since contracts do not have to be negotiated with every worker individually (see also Prendergast 1999). In addition, a concern for fairness has been a main argument invoked to justify equal wages. It has been argued that differential pay of co-workers is considered unfair by workers, causes resentment and envy within the workforce, and ultimately lower performance (Pfeffer and Langton 1993, Bewley 1999). Equality is also often referred to in employer-union bargaining as being a cornerstone of a fair wage scheme.

However, it could be that wage equality hampers the effectiveness of gift exchange. This is likely to be the case if agents do not primarily consider wage *equality* as fair, but rather care about horizontal *equity* (Adams 1963). In a work environment, the equity principle (or "equity norm") demands that a person who exerts higher effort should receive a higher wage compared to his co-worker. Only when performance of co-workers is the same, equity and equality coincide.

In this chapter, we study the relative importance of these fairness principles by analyzing their efficiency implications when contract enforcement relies on gift exchange. We do so in a simple and parsimonious laboratory experiment in which one principal interacts with two agents. In a first stage the agents exert costly effort. After observing their efforts, the principal pays them a wage. In the main treatment he can choose the level of the wage but he is obliged to pay the same wage to both agents (*equal wage treatment* or EWT). In the control treatment, the principal can wage discriminate between the two agents (*individual wage treatment* or IWT). In both treatments, neither efforts nor wages are contractible.

The main findings of the experiment are as follows. First, performance differs substantially between the EWT and the IWT: agents who are paid equal wages exert significantly lower efforts than agents who are paid individually. Effort levels are nearly twice as high under individual wages. In addition, efforts decline over time when equal wages are paid. Second, this strong treatment effect cannot be explained by differences in monetary incentives. The actual wage choices of principals imply that providing high effort levels is profitable for agents in both treatments. From a purely monetary viewpoint agents' behavior in both treatments should thus be similar. Third, we show that the frequent violation of the equity norm in the equal wage treatment can explain the effort differences between the treatments. In both treatments, agents who exert a higher effort and earn a lower payoff than their co-worker strongly decrease their effort in the next period. This pattern is very similar in both treatments. However, the norm of equity is violated much more frequently under equal wages. Principals in the IWT understand these mechanisms quite well. When efforts differ they do pay different wages, rewarding the harder-working agent with a higher payoff in most cases.

Our results suggest a psychological rationale for using individual wages. Subjects perceive equal wages for unequal performance as unfair and reduce their effort subsequently. The traditional literature on incentive provision in groups comes to a similar conclusion though for a different reason. It is usually argued that the inefficiency of equal wages stems from the fact that marginal products and wages are not aligned. This can lead to free-riding among selfish agents (e.g., Holmström 1982, Erev et al. 1993). We enlarge the scope of this critical view on wage equality: interestingly, in our setup it is precisely the presence of fair-minded agents and not their absence that calls for the use of individual rewards.

Regarding compensation practice in firms, our findings highlight the importance of taking the concerns for co-workers' wages into account. However, doing so by paying equal wages to a group of agents may actually do more harm than good. As soon as agents differ in their performance, equal wages which seem to be a fair institution at first sight might be considered very unfair. While the discouraging effect of equal wages on hard-working agents has long been informally discussed (e.g., Milgrom and Roberts 1992, p. 418f) this chapter provides controlled evidence in favor of this intuition. Moreover, it suggests that it is the violation of the norm of equity that causes the discouragement and low performance.

The results of Chapter 2—in particular, the performance of agents under individual wages—show how powerful gift exchange can be in eliciting work effort: although explicit contract enforcement is not feasible, 80% of the possible efficiency gains are realized. However, while gift exchange and other forms of implicit contracts might be successful in eliciting high work efforts, they could have less desirable consequences for other labor market outcomes. In particular, some instruments of implicit contract enforcement could give rise to involuntary unemployment. For instance, when employment relations are based on fair wages and gift exchange, it could be optimal for firms to dismiss workers rather than cut wages in times of economic downturns (Bewley 1999). More generally, firms might prefer to ration jobs because the requirement to pay fair wages could render less productive jobs unprofitable. Finally, the presence of unemployed workers could itself be necessary to motivate employed workers to work hard (Shapiro and Stiglitz 1984, MacLeod and Malcomson 1989).

In Chapter 3, we empirically analyze the relationship between contract enforcement and the emergence of unemployment. In particular, we are interested in whether the absence of third party contract enforcement has a direct impact on the level of unemployment. Moreover, we study how different instruments of implicit contracting—such as reputation in relational contracts, gift exchange, and an implicit threat of dismissal—interact in determining labor market outcomes. Do firms, for instance, ration jobs but share production rents generously when effort is not verifiable? Do they engage in repeated, long-term work relationships? How does unemployment influence workers' effort provision?

We address these questions in an experimental labor market where we exogenously vary the verifiability of work effort. In the market, firms and workers interact during multiple market periods. In every period, each firm can hire up to two workers. In addition to the analysis in the previous chapters, we allow firms and workers to trade repeatedly. This allows the trading parties to build up long-term work relationships and make use of relational contracts as an additional means of providing incentives. Our two experimental treatments only differ in the degree to which work effort of employed workers is explicitly enforced. In our control treatment, concluded contracts are exogenously enforced, i.e., a worker's effort has to be equal to the desired effort level stipulated in the employment contract. By contrast, effort in our main treatment is not verifiable, and firm and workers therefore have to rely on implicit incentives.

Our experiment yields the following findings. Unemployment is much higher in the treatment where effort is not verifiable. More importantly, unemployment in this treatment is involuntary, being caused by the firms' employment and contracting policy. Firms pay "fair" wages but offer fewer vacancies than possible and technologically efficient. By contrast, wages are close to the market clearing level when contracts are explicitly enforced. Firms do not ration jobs, and unemployment is very low. Moreover, unemployment in this treatment is mostly voluntary, being caused by workers who do not accept existing contract offers.

The firms' employment policy in the treatment without explicit contract enforcement, however, again succeeds in eliciting high efforts from the employed workers. One reason why this is the case is that many firms build up repeated relationships with specific workers. These long-term work relationships are characterized by high wages and high effort levels, thus yielding positive rents both for firms and workers. Providing high efforts in response to high wages is profitable for the workers because of two reasons. First, firms do not rehire workers who shirked in previous periods. At the same time, a job loss entails considerable costs for the workers: due to the high level of unemployment, the job acquisition rate for unemployed workers is very low. Our findings therefore lend support to theories of efficiency wages based on gift exchange (e.g. Akerlof 1982) as well as models of unemployment as a worker disciplining device (e.g. Shapiro and Stiglitz 1984).

In Chapter 4, we analyze a different source of performance incentives for employees in many occupations—the prospect of being promoted to better paid positions at higher organizational levels. According to Lazear and Gibbs (2008), promotions are probably the most important source of extrinsic motivation for middle managers in most firms. Using promotions as an instrument for incentive provision is especially important when other, more direct means of pay-for-performance are absent, e.g., because it is not feasible to measure individual performance objectively. In the multi-level hierarchies of most modern organizations, employees promoted at one level have the chance to compete for further promotions, giving rise to "elimination" tournaments with multiple stages. In such elimination tournaments, the incentive effect of a promotion does not only stem from the immediate wage increase but also from the option value of further promotion chances.

Although the importance of multi-stage elimination tournaments is undisputed, stringent empirical tests of their incentive effects have so far been scarce. In Chapter 4 we experimentally study the implications of multi-stage tournaments for employee performance in comparison to single promotion decisions using controlled laboratory techniques. In addition, we address the question how different wage structures, i.e., the size of wage jumps between different levels of a hierarchy affect performance.

We study these questions by comparing three treatments. Our main treatment is a two-stage tournament (TS) in which four subjects compete for being promoted to an intermediate level. In a second stage the two promoted subjects compete again for the top position. The promotion decision depends on costly effort and an individual noise term. We compare this treatment to a one-stage tournament (OS) in which four subjects compete once for two top positions. The wage for the promoted subjects is chosen such that the one-stage tournament is strategically equivalent to the first stage of TS, i.e., equilibrium efforts are the same in both treatments. Comparing OS and TS thus allows testing whether subjects take the option value in multi-stage tournaments into account and evaluate it correctly. In a third treatment (TSC), we analyze whether a more convex wage profile leads to reduced first stage efforts and increased efforts in the second stage of the elimination tournament as predicted by theory.

Our findings can be summarized as follows: First, confirming previous evidence, the predictions of tournament theory are remarkably close to average behavior in our one-stage treatment. Second, behavior in the TS treatment indicates that subjects take into account the option value of future promotion possibilities when deciding on their work effort in multi-stage tournaments. However, we also observe important departures from theoretical predictions in this tournament. Behavior in the first stage of TS differs strongly both from the one-stage treatment and from theoretical predictions. Subjects exert significantly higher effort in the first stage of the two-stage tournament. This excess exertion of effort is, however, not observed in the second stage of the two-stage tournament. Finally, the results of the TSC treatment suggest that excess effort in the first stage of two-stage tournaments is a robust finding. Subjects hardly react to the change in the wage structure, implying that first-stage excess effort is even higher in TSC.

Our results enhance the understanding of incentive effects in complex multi-stage promotion tournaments. They indicate that the mechanisms of incentive provision in multi-stage tournaments largely operate as suggested by theory. People do not only respond to the immediate increase in wages but also seem to be motivated by the option value entailed in future promotion possibilities. Our findings also provide insights with regard to the question whether one-stage tournaments are behaviorally equivalent to multi-stage designs. Adding one or more stages seems to make a fundamental difference, as people tend to exert excess effort in early stages of the tournament. This shows that one cannot necessarily draw inferences from simple one-stage setups to more complex tournaments.

The finding that people tend to exert excess effort in early stages of a multi-stage competition also has potentially wide-ranging organizational implications. It suggests that multi-stage elimination tournaments can be a particularly effective form of providing incentives because they elicit unusually high work efforts at the lower levels of a hierarchy. Ultimately, this could help to explain why firms rely so heavily on promotion based incentive schemes even if more direct means of performance assessment and compensation are available.[1]

[1]Chapter 1 is based on joint work with Thomas Dohmen and Matthias Wibral (Altmann et al. 2008). Chapter 2 was developed jointly with Johannes Abeler, Sebastian Kube, and Matthias Wibral. An earlier version of this chapter was circulated under the title "Reciprocity and Payment Schemes: When Equality is Unfair". Chapter 3 owes to the collaboration with Armin Falk and David Huffman. Chapter 4 is joint work with Armin Falk and Matthias Wibral and was circulated under the title "Promotions and Incentives: The Case of Multi-Stage Elimination Tournaments".

Chapter 1

Do the Reciprocal Trust Less?

1.1 Introduction

By now there seems to be broad agreement that trust and reciprocity are conducive to economic performance and efficiency (e.g., Knack and Keefer 1997). Mutual trust between trading parties facilitates the realization of gains from trade, for instance by reducing contracting costs. Reciprocity can also enhance performance in many areas of economic life, for example by mitigating moral hazard problems in labor relations (Fehr et al. 1997). In order to better understand the economic implications of reciprocity, several formal models of social preferences have been developed (e.g., Fehr and Schmidt 1999, Falk and Fischbacher 2006).

In spite of the importance of trust and reciprocity surprisingly little is known about their relationship on an intrapersonal level. In other words, do reciprocal persons trust more or less than selfish ones? In this chapter we address precisely this question with the help of a controlled laboratory experiment, employing a variant of the trust game that allows us to measure both variables for each individual.[1]

[1] Several studies (e.g., Cox 2004, Ashraf et al. 2006) have analyzed behavior across games in order to disentangle subjects' unconditional kindness or altruism from trust and reciprocity, but these studies do not look at the direct link between a person's (own) reciprocity and trust. In addition, it is not clear to what extent inferences can be made from behavior in non-strategic environments (e.g., the dictator game) to players' motives in strategic interactions. See Fehr and Schmidt (2006) for a discussion of this point.

We find a strong and positive relationship between a person's reciprocity and her trusting behavior. Reciprocal players exhibit much higher levels of trust than more selfish ones, even when personal characteristics and preferences such as gender or risk attitudes are controlled for. This finding is also interesting from a theoretical perspective because theories of social preferences typically assume—at least implicitly—a connection between trust and reciprocity. In particular, the observed positive relation between the two raises important questions about theories which predict that *ceteris paribus* "fairer" players trust *less*.

The remainder of this chapter is organized as follows. The next section describes the design of our experiment, Section 1.3 presents the empirical results. Section 1.4 concludes by discussing the implications of our findings for modelling social preferences.

1.2 Experimental Design

In our experiment, subjects were anonymously matched in pairs and played a modified version of the trust game (Berg et al. 1995). Both players received an endowment of 120 points. The first mover (the *sender*) could send any amount $t \in \{0, 20, 40, 60, 80, 100, 120\}$ to the second mover (the *receiver*). The amount sent was tripled by the experimenter. Then, the second mover could send back any amount between zero and 480 points. The crucial feature that distinguishes our design from the original version of the trust game is the use of the strategy method to elicit each subject's trust and reciprocal inclination. In our experiment, subjects made decisions both in the role of the sender and the receiver. In the role of the receiver subjects had to decide how much to send back for any possible amount received. This procedure allows us to measure both the level of trust and the level of reciprocity for each subject in the same strategic environment.[2]

To give subjects the monetary incentives to take all decisions seriously while

[2]Other studies have employed the strategy method in trust games in which subjects play only one role (e.g., Bellemare and Kröger 2007, Falk and Zehnder 2007). Burks et al. (2003) have subjects play both roles but do not use the strategy method.

at the same time avoiding potential confounds if subjects interact repeatedly in different roles, we employed the following incentive-compatible procedure. After all decisions had been made, a random mechanism determined which player of a given pair actually had the role of the sender and which player had the receiver role. Then, players' decisions were implemented and subjects were paid accordingly.

The experiment was programmed with the software z-Tree (Fischbacher 2007) and conducted at the BonnEconLab. Twenty subjects participated in each of the 12 sessions that we ran so that we observe the choices of 240 different subjects. The trust game was part of a sequence of tasks (see Dohmen and Falk 2006 for a detailed description). Before subjects played the trust game they had to solve math problems under different monetary incentives.[3]

After the trust game, we elicited subjects' risk attitudes using a series of 15 choices between a safe payment and a lottery. The lottery was the same across choices (400 points or 0 points, each with probability 0.5) while the safe option increased from 25 points to 375 points in increments of 25. If subjects have monotonous preferences, they prefer the lottery up to a certain level of the safe option, and then switch to preferring the safe option in all subsequent choices. After a subject had made decisions for all 15 choices, it was randomly determined which choice became relevant for payment.[4] Together, the trust game and the lottery choice task lasted about 20-25 minutes and subjects earned 6.87 Euro on average.

1.3 Results

We measure *trust* by the amount that a subject sends as a first mover. Our measure of *reciprocity* (also denoted "r") is derived as follows: for each subject, we used the

[3] In 4 of the 12 sessions subjects worked under purely individual incentives (fixed wages and piece rates). In the remaining 8 sessions they could select into an incentive scheme (team or tournament) which involved anonymous interaction with another player. All subjects were randomly rematched in the trust game. In view of our results we are confident that neither solving math problems nor the different incentive schemes systematically affect behavior in the trust game (see below).

[4] The experimental instructions for the trust game can be found in Appendix A.1. Except for the payoff parameters, the lottery procedure was identical to the one described in Appendix D.7.

decisions as a second mover and ran an OLS-regression of the amounts sent back on the (hypothetical) amounts sent by the opponent, forcing the slope through the origin. The slope coefficient gives us a measure of a subject's willingness to reward kind actions of an opponent by own kind behavior, i.e., positive reciprocity. If a receiver, for example, always matches his final payoff with that of the sender, his reciprocity coefficient is $r = 2$.

In order to graphically present our main result, we classify subjects according to their behavior as second movers. We call subjects with a reciprocity parameter $r > 1$ "reciprocal", and subjects with a slope parameter $r = 0$ "selfish". Reciprocal types leave their opponent with a positive return to trust, sending back more than the amount sent to them by the sender. 64.6% of our subjects fall into this category. Selfish types, who make up 12.5% of subjects, never send back anything, irrespective of the first mover's behavior. The remaining 22.9% of subjects whose slope parameter is positive, but small ($r \leq 1$) are categorized as "intermediate" types.

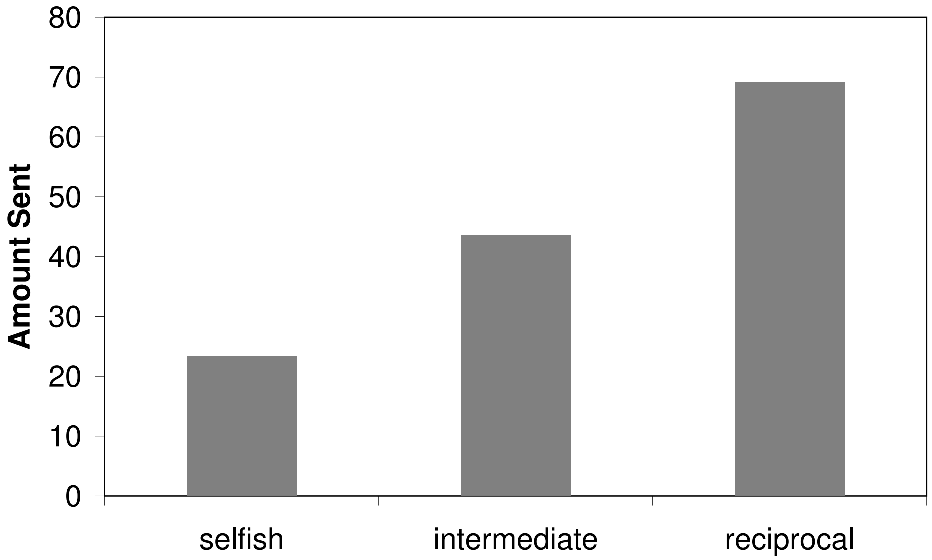

Figure 1.1: *Average amount sent by selfish, intermediate, and reciprocal players.*

Figure 1.1 plots the amount that the three types of subjects send on average in the trust game. Reciprocal types clearly send most (69.2 points on average), and selfish types send least (23.3 points). Subjects in the intermediate category send 43.6 points. Pairwise Mann-Whitney-U-tests indicate that all differences between the groups are highly statistically significant ($p < 0.01$ for all pairwise tests). Subjects who always "split the pie equally" as a responder (i.e., subjects with $r = 2$) trust most (83.5 points). The result depicted in Figure 1.1 is robust to different classifications of types using a finer "grid".

OLS-regressions of individuals' trust, measured by the amount sent in the trust game, on their reciprocal inclination, measured by the slope parameter r described above, confirm that reciprocal individuals trust more: an increase of one unit in the reciprocity measure is associated with sending 16.1 points more in the trust game (see Column (1) of Table 1.1).[5]

	Dependent variable: Amount sent		
	(1)	(2)	(3)
Reciprocity	16.091***	17.234***	17.761***
	(2.235)	(2.199)	(2.316)
1 if male		16.301***	15.167***
		(4.401)	(4.688)
Certainty equivalent			0.100**
			(0.044)
Constant	34.081***	24.329***	3.665
	(3.957)	(4.668)	(10.151)
R^2 adj.	0.175	0.217	0.230
Observations	240	240	221

Table 1.1: *Trust-Regressions. OLS estimates (standard errors in parentheses).* "Certainty Equivalent" *indicates the switch from the risky lottery to the safe option (=0 if subject is strongly risk averse, ..., =400 if subject is strongly risk loving). Significance at the 10%, 5% and 1% level is denoted by *, **, and ***, respectively.*

This key result is also robust to controlling for gender (cf. Column (2) of Table 1.1) and subjects' risk attitudes (cf. Column (3) in Table 1.1).[6] The positive influence

[5] This result still holds if we restrict the sample to the 4 sessions in which there was no interaction between subjects in the tasks preceding the trust game. In addition, we find the significant positive correlation between trust and reciprocity in the remaining 8 sessions irrespective of the chosen incentive scheme.

[6] The certainty equivalent cannot be determined unambiguously for 19 subjects because they

of reciprocity on trust is highly significant and also quantitatively very similar in all specifications. The effects of gender and risk attitudes are consistent with the findings in the literature. Men trust more than women (cf. Bohnet and Zeckhauser 2004), sending about 15 points more than female participants in our sample. Our results also confirm the importance of risk attitudes for trusting behavior: subjects who are more willing to take risks send significantly more (cf. Dohmen et al. 2006).

1.4 Discussion and Concluding Remarks

The strong, positive relationship between a person's reciprocal inclination and her trusting behavior has important implications for the evaluation and advancement of theories that incorporate social preferences. Some of the most prominent models (e.g., Fehr and Schmidt 1999, Bolton and Ockenfels 2000, Falk and Fischbacher 2006) predict that individuals who are more reciprocal (or inequity averse) ceteris paribus trust *less* than others in the trust game. The intuition for this result is that a selfish sender just suffers from the loss of her investment if the receiver sends back too little, whereas a fair-minded sender experiences additional disutility because his trust has been exploited.[7]

Our results show, however, that the relationship between trust and reciprocity may be more complex than captured by most models. The finding that people trust more the more reciprocal they are allows at least two different preliminary interpretations—one based on norm adherence and the other on systematic differences in beliefs. The idea of the former is that some people value adherence to a certain moral norm in itself. If these people follow a norm that, e.g., dictates cooperative behavior in either role, this could account for our main finding. Such norm-guided behavior could also help to explain why some senders in trust games send positive amounts despite expecting to get back less than they send (cf. Dufwen-

switched more than once between the safe option and the lottery. These subjects were excluded from the regression in Column (3). Including them with the lowest or highest switching point from the lottery to the safe option does not change the results.

[7]Along these lines, Fehr et al. (2007) have argued that "fairness preferences inhibit trusting behavior because trust typically involves a risk of being cheated."

berg and Gneezy 2000, Ashraf et al. 2006).

A different interpretation is that fair and selfish types have fundamentally different beliefs regarding the behavior of others. Such differences in beliefs might be the result of a "false consensus effect" (Kelley and Stahelski 1970). As an extreme example, assume that reciprocal players expect all others to behave reciprocally, and that a selfish subject expects all others to be selfish as well. In this case reciprocal types will send positive amounts and expect a positive return, while selfish types will never send anything since they expect that the receiver will not send anything back. Such systematic differences in beliefs would have interesting implications for the modelling of social preferences as they require giving up the widely used common-prior assumption. They potentially also have important practical implications as they could lead different types of players to select into different institutional settings. This could help to explain why environments with different degrees of exogenous enforcement coexist, e.g., in the labor market. Which of the two interpretations is more relevant cannot be answered with our data but remains an important question for future research.

Chapter 2

When Equality is Unfair

"To treat people fairly you have to treat people differently."
Roy Roberts, at that time VP of General Motors[1]

2.1 Introduction

In recent years, a vast body of literature has stressed the importance of gift exchange for mitigating moral-hazard problems of incomplete contracts: since many agents repay a gift in the form of higher wages by providing higher efforts, effort can be elicited under contractual incompleteness even in one-shot situations where no future gains can be expected (e.g., Akerlof 1982, Fehr et al. 1997, Maximiano et al. 2007). The potential of gift exchange as a contract enforcement device, however, is likely to depend on the institutions that shape the employment relation, above all the mode of payment. Yet little is known about the interaction of different payment modes with gift exchange. Exploring this interaction is crucial in order to understand under which conditions the efficiency-enhancing effects of gift exchange develop their full power. A key question in this context is how to treat agents relative to each other as this affects the perceived fairness of a pay scheme. In this chapter, we study this question by focusing on two important fairness principles: horizontal equality and equity.

[1]Quoted in Baker et al. (1988).

On the one hand, it has been argued that *horizontal equality* is crucial for a wage scheme to be considered as fair. Differential pay of co-workers could cause resentment and envy within the workforce, and ultimately lower performance (e.g., Pfeffer and Langton 1993, Bewley 1999). Wage equality is also often referred to in employer-union bargaining as being a cornerstone of a fair wage scheme and is one of the most prevalent payment modes (see, e.g., Medoff and Abraham 1980, Baker et al. 1988). If workers care foremost about equality, a wage scheme that guarantees equal wages for co-workers should lead to an efficiency-enhancing gift exchange relation. On the other hand, the importance of the *equity principle* has long been discussed in social psychology, personnel management, and economics (e.g., Homans 1961, Fehr and Schmidt 1999, Konow 2003). In a work environment, the equity principle (or "equity norm") demands that a person who exerts higher effort should receive a higher wage compared to his co-worker. Only when performance of co-workers is the same, equity and equality coincide. However, in real-life work relations this is likely to be the exception rather than the rule. Whenever workers differ in their performance, horizontal wage equality violates the equity principle since a higher effort is not rewarded with a higher wage. In other words, if equity is important, the often-heard slogan "equal pay for equal work" implies "unequal pay for unequal work".[2]

Ideally, our research question would be examined in work environments that differ only with respect to the payment mode. To come close to this ideal world, we introduce a simple and parsimonious laboratory experiment that allows us to analyze the interaction between the institution of wage equality and gift exchange. In the experiment, one principal is matched with two agents. In a first stage the agents exert costly effort. After observing their efforts, the principal pays them a wage. In the main treatment he can choose the level of the wage but he is obliged to pay the same wage to both agents (*equal wage treatment* or EWT). In

[2]Lazear (1989) neatly summarizes this discussion (p. 561): "It is common for both management and worker groups such as labor unions to express a desire for homogeneous wage treatment. The desire for similar treatment is frequently articulated as an attempt to preserve worker unity, to maintain good morale, and to create a cooperative work environment. But it is far from obvious that pay equality has these effects."

the control treatment, the principal can wage discriminate between the two agents (*individual wage treatment* or IWT). In both treatments, neither efforts nor wages are contractible. Note that principals in the individual wage treatment are free to pay the same wage to both agents, i.e., the EWT is a special case of the IWT. If agents care foremost about wage equality, there should thus be no treatment difference; if equity considerations are more important, we should find that the EWT elicits lower effort levels than the IWT.

The main findings of the experiment are as follows. First, performance differs substantially between the EWT and the IWT: agents who are paid equal wages exert significantly lower efforts than agents who are paid individually. Effort levels are nearly twice as high under individual wages and efforts decline over time when equal wages are paid. Second, this strong treatment effect cannot be explained by differences in monetary incentives. The actual wage choices of principals imply that providing high effort levels is profitable for agents in both treatments. From a purely monetary viewpoint agents' behavior in both treatments should thus be similar. Third, we show that the frequent violation of the equity principle in the equal wage treatment can explain the effort differences between the treatments. In both treatments, agents who exert a higher effort and earn a lower payoff than their co-worker strongly decrease their effort in the next period. However, the norm of equity is violated much more frequently under equal wages. Principals in the IWT understand the mechanisms of equity quite well. When efforts differ they do pay different wages, rewarding the harder-working agent with a higher payoff in most cases. Agents' reactions cause completely different dynamics in the two main treatments. Under equal wages, initially hard-working agents get discouraged and reduce their effort to the level of their low-performing co-workers. By contrast, in the individual wage treatment the high performers keep exerting high efforts while the low performers change their behavior and strongly increase their effort levels.

Note that principals in the IWT can set two wages instead of one in the EWT. This opens the possibility that agents attribute a different degree of intentionality to principals' wage choices. It could be that this additional moment of discretion has a direct impact on the treatment difference. To rule out this potential confound, we conduct an additional control treatment where principals can again set only one wage as in the EWT. The second wage is set *exogenously* such that the equity principle is always fulfilled. Effort levels in the control treatment are similar to those of the IWT and much higher compared to the EWT. This confirms that the difference between our two main treatments is indeed driven by agents' desire for wages that are in line with the equity principle.

Our results suggest a psychological rationale for using individual wages. Subjects perceive equal wages for unequal performance as unfair and reduce their effort subsequently. The traditional literature on incentive provision in groups comes to a similar conclusion though for a different reason. It is usually argued that the inefficiency of equal wages stems from the fact that marginal products and wages are not aligned. This can lead to free-riding among selfish agents (e.g., Holmström 1982, Erev et al. 1993). We enlarge the scope of this critical view on wage equality: interestingly, in our setup it is precisely the presence of fair-minded agents and not their absence that calls for the use of individual rewards.

An earlier literature in social psychology also studies the consequences of equity in social exchanges (Homans 1961, Adams 1963, Adams 1965, Andrews 1967). In his influential equity theory, Adams (1965) operationalizes the general equity principle in an "equity formula", which states that the ratio of outcomes to inputs should be the same for every individual.[3] If this is not the case an individual experiences distress and seeks to reestablish equity. Our study complements this literature in several ways. As Mowday (1991) notes, interpreting the existing empirical evidence can often be difficult because important aspects such as the costs of effort or the relevant reference group are ambiguous. Our economic laboratory experiment offers a high level of control over these aspects. In addition, violations of the equity norm arise from the interaction of principals and agents in our study whereas they are induced by the experimenter in most earlier experiments, e.g., by making subjects believe they are over- or underqualified for the job (e.g., Adams 1963 or Lawler 1967).

[3]The idea of proportionality dates back to at least Aristotle's Nicomachean Ethics.

Since agents in our experiment compare their payoff with the payoff of their co-worker, our results also inform the literature analyzing the influence of relative income on satisfaction and performance. It has been shown that relative income affects people's well-being (e.g., Clark and Oswald 1996, Easterlin 2001, Fließbach et al. 2007). However, it is less clear how this influences performance, i.e., whether low relative income leads to frustration and reduced performance (as in Clark et al. 2006 and Torgler et al. 2006) or to an increase in performance due to a "positional arms race" (Neumark and Postlewaite 1998, Layard 2005, Bowles and Park 2005). The controlled laboratory environment of our experiment allows us to reconcile these differing views. Our results indicate that the comparison process goes beyond a one-dimensional comparison of income and also includes a comparison of effort. In particular, they suggest that receiving a lower income while exerting a *higher* effort leads to reduced performance as this conflicts with the equity principle. By contrast, a lower income that is generated by a *lower* effort leads to a (small) increase in performance.

There are only a few experimental studies that analyze the interaction of payment modes and social preferences (e.g., Bandiera et al. 2005, Fehr et al. 2007, Falk, Huffman and MacLeod 2008). Most closely related to this chapter is the work of Charness and Kuhn (2007). Here, one principal is matched with two agents differing in productivity; like in our study, wages and efforts are not contractible. In contrast to our results, they find that co-workers' wages do *not* matter much for agents' decisions. However, their design differs from ours in several important points. While Charness and Kuhn focus on heterogeneity in productivity, we look at the effect of actual output differences between agents. Furthermore, we allow for richer comparisons between the agents, as in their design agents are not aware of the magnitude and direction of the productivity differences. The different results underline the importance of information for determining the reference group: Charness and Kuhn's results rather apply to groups of workers that are loosely related and know little about each other, while our focus is on close co-workers who have a good understanding about their peers' abilities and efforts.

Regarding compensation practice in firms, our findings highlight the importance of taking the concerns for co-workers' wages into account. However, doing so by paying equal wages to a group of agents may actually do more harm than good. As soon as agents differ in their performance, equal wages which seem to be a fair institution at first sight might be considered very unfair. While the discouraging effect of equal wages on hard-working agents has long been informally discussed (e.g., Milgrom and Roberts 1992, p. 418f) we provide controlled evidence in favor of this intuition. Moreover, our findings suggest that it is the violation of the norm of equity that causes the discouragement and low performance. Our results should not be interpreted as arguments against wage equality in general. They rather point to limits of equal wages.[4] Wage equality is potentially a good choice in occupations where, e.g., due to technological reasons, workers' performance differs only slightly or where performance differences are due to random influences. In addition, the transparency of co-workers' work efforts and wages might have an influence on the optimal choice of the pay scheme.

The remainder of this chapter is structured as follows. In the next section we describe the experimental design and discuss theoretical predictions. In Section 2.3 we present and discuss our results and Section 2.4 concludes.

2.2 Experimental Setup

2.2.1 Design and Procedures

In the experiment, one principal is matched with two agents. The subjects play a two-stage game. In the first stage, agents decide simultaneously and independently how much effort they want to provide. Exerting effort is costly for the agents. Effort choices range from 1 to 10 and are associated with a convex cost function displayed in Table 2.1. The principal reaps the benefits of production: every unit of effort increases his payoff by 10.

[4]Independent of equity-equality trade-offs, equal wages might be beneficial for the principal because they could increase peer monitoring (Knez and Simester 2001) and lower transaction costs since contracts do not have to be negotiated with every worker individually (e.g., Prendergast 1999).

Effort level e_i	1	2	3	4	5	6	7	8	9	10
Cost of effort $c(e_i)$	0	1	2	4	6	8	10	13	16	20

Table 2.1: *Cost of effort.*

In the second stage, after observing the effort decisions of his agents, the principal decides on wages for the two agents. The wages have to be between 0 and 100. Neither efforts nor wages are contractible. The only difference between treatments is the mode of payment. In our main treatment the principal can only choose one wage w that is paid to each of the agents (*equal wage treatment* or EWT). In the control treatment he can discriminate between the two agents by choosing wages w_1 and w_2 for agent 1 and 2, respectively (*individual wage treatment* or IWT). The EWT is thus a special case of the IWT. At the end of each period, the two agents and the principal are informed about efforts, wage(s), and the resulting payoffs for all three players. The payoff functions for the players are summarized in Table 2.2.

Treatment	EWT	IWT
Payoff Principal	$\pi_P = 10(e_1 + e_2) - 2w$	$\pi_P = 10(e_1 + e_2) - (w_1 + w_2)$
Payoff Agent i	$\pi_{A_i} = w - c(e_i)$	$\pi_{A_i} = w_i - c(e_i)$

Table 2.2: *Payoffs of players.*

This game is played for twelve periods. We implemented a stranger design to abstract from confounding reputation effects, i.e., at the beginning of each period principals and agents were rematched anonymously and randomly within a matching group. A matching group consisted of three principals and six agents. The subjects kept their roles throughout the entire experiment. After the last period, subjects answered a short post-experimental questionnaire. The experiment was conducted in a labor market framing, i.e., principals were called "employers" and agents were called "employees".[5]

[5] An English translation of the instructions can be found in Appendix B.2.

Our setup is related to the gift-exchange game introduced by Fehr et al. (1993) but differs in two important ways. First, in our experiment agents move first while in Fehr et al.'s setup the principal moves first. Our move order allows the principal to base his wage decision on the actually exerted effort. More importantly, a principal in our experiment is matched with two agents instead of one. This is an essential prerequisite to analyze the interaction between gift exchange and payment modes. It allows us to study the impact of relative wages on the perceived fairness of the wage scheme and agents' behavior.

All participants started the experiment with an initial endowment of 400 points that also served as their show-up fee. Points earned were converted at an exchange rate of 0.01 Euro/point. The experiment was conducted at the BonnEconLab at the University of Bonn in April 2005 using z-Tree (Fischbacher 2007). For each treatment, we ran four sessions with a total of 8 matching groups (144 participants). The experiment lasted approximately 70 minutes. On average subjects earned 8.30 Euro.

2.2.2 Behavioral Predictions

Efficiency is determined by agents' effort choices. It is maximized if both agents exert the highest possible effort of 10. However, if all players are rational and selfish the principal will not pay anything to the agents since wage payments only reduce his monetary payoff. Anticipating this, both agents will provide the minimal effort of one in the first stage. The finite repetition of the game in randomly rematched groups does not change this prediction. This subgame perfect equilibrium is the same for both payment modes. If all players were selfish we should therefore expect no difference between treatments.

By contrast, in laboratory experiments studying labor relations with incomplete contracts, one typically observes that efforts and wages exceed the smallest possible value. Moreover, wages and efforts are positively correlated (e.g., Fehr and Gächter 2000). These findings illustrate the potential of reciprocal gift exchange in enforcing incomplete contracts, as postulated in Akerlof and Yellen's fair wage-effort hypothesis (Akerlof and Yellen 1990). A fundamental prerequisite for the functioning of gift-exchange relations is that workers perceive their wage as fair. The fairness of a wage payment, however, may not only be evaluated in absolute terms, but also *relative* to the wages of other members in a worker's reference group.[6] This is not important for the special case of bilateral gift-exchange relationships where only one agent interacts with one principal (e.g., Fehr et al. 1997). However, horizontal fairness considerations potentially play a crucial role in our setup where workers can compare to co-workers.

How do the behavioral predictions depend on which horizontal fairness principle is most important? If agents in the experiment care foremost about *wage equality*, the EWT—which guarantees equal wages by design—should lead to efficient gift exchange between firms and workers. Additionally, we should expect no behavioral differences between treatments since firms in the IWT can pay their workers equal wages, too. Given that firms in the IWT recognize workers' desire for equal treatment, they will decide to do so. Thus, the wage-effort relationship and average effort levels should not differ across treatments. If some firms do nevertheless wage discriminate between workers, the IWT should lead to less efficient outcomes than the EWT.

By contrast, if workers consider *equity* to be more important than equality, we should expect differences in behavior between treatments. The equity principle demands that a person who exerts higher effort than his co-worker should receive a higher wage and payoff. Our experimental treatments differ in the extent to which the equity principle can be fulfilled by principals. Under the equal wage institution, the equity norm is violated whenever agents differ in their performance. Since both workers receive the same wage but have to bear the cost of effort provision, the worker who exerts more effort receives a lower monetary payoff. Under individual wages, principals' behavior determines endogenously whether the equity norm is violated or not. By differentiating wages in accordance to effort differences, principals

[6]Potentially many variables influence a worker's fairness perception of his wage, e.g., the unemployment rate, unemployment benefits, the prevailing market wage, etc. (see Akerlof 1982, Akerlof and Yellen 1990). These factors are ruled out by our experimental design, allowing us to isolate the influence of co-workers' wages on fairness perception and effort provision.

can adhere to the norm. If we assume that at least some principals do so, we expect to see less norm violations in IWT than in EWT.

What are the behavioral consequences of such differences in norm fulfillment? Agents who value equitable treatment should suffer from norm violations, feel dissatisfied and subsequently try to restore equity by adjusting their behavior. Equity theory proposes several possible reactions of agents after norm violations, such as altering own or others' efforts or payoffs, changing one's reference group or quitting the relationship (see Adams 1965). The virtue of our experimental design is that we can clearly identify agents' reactions, because the only variable that an agent can change after experiencing a norm violation is his work effort. An agent who faces a disadvantageous norm violation (i.e., relative underpayment) should lower his effort in the following period. An agent who experiences an advantageous norm violation (i.e., relative overpayment) should increase his effort. Note that a norm violation always includes one agent facing a disadvantageous violation and one agent facing an advantageous violation. Dissatisfaction and the resulting strength of reactions, however, is likely to depend on the direction of the norm violation. Previous evidence suggests that the decrease of effort after a disadvantageous norm violation will be stronger than the increase of effort after an advantageous violation (Loewenstein et al. 1989, Mowday 1991, Thöni and Gächter 2008). Consequently, a violation of the equity norm should lead to an *overall decrease* of efforts in the subsequent period.

If workers care about equitable payment in the sense of the postulated equity norm, aggregate effort in the EWT should thus be lower compared to the IWT since we expect to observe less norm violations in the latter.

2.3 Results

In this section we present the results of the experiment and discuss possible explanations for the observed behavior. We first analyze efficiency implications of the two payment schemes by comparing the effort choices of agents. We then demonstrate that the difference in agents' performance obtains even though monetary incentives—implied by principals' wage setting—should lead to similar effort choices in both treatments. Subsequently, we show that workers' behavior is strongly affected by the equity principle, which is more frequently violated in the EWT. Finally, we report the results of an additional control experiment. They demonstrate that the higher efficiency of the IWT is not driven by the fact that principals can set two wages instead of one (as in the EWT) but by the fact that principals set wages that are in line with the equity principle.

2.3.1 Effort Choices and Efficiency

Figure 2.1 shows the development of average efforts over time. Under equal wages, efforts are lower already in the first period (Mann-Whitney test: $p = 0.03$)[7] and decrease over time. Efforts under individual wages stay constant (Wilcoxon test for periods 1–6 against 7–12: IWT, $p = 0.56$; EWT, $p < 0.01$). This results in a strong overall treatment difference: average efforts are almost twice as high in the IWT compared to the EWT (8.21 vs. 4.40; Mann-Whitney test: $p < 0.01$). The treatment difference is also present when individual matching groups are considered: the highest average effort of an EWT matching group (5.88) is still lower than the lowest average effort of an IWT matching group (7.47).

The difference in agents' behavior can also be seen in the histogram of effort choices (Figure 2.2). In the individual wage treatment agents choose the maximum effort of 10 in 49% of the cases, 84% of the choices are higher than 6. Under equal wages, agents choose an effort higher than 6 in only 26% of all cases. The effort decisions are more spread out in the EWT, the minimal effort of 1 being the modal choice with 24% of the choices. Since higher efforts increase production and since the marginal product of effort always exceeds its marginal cost, the differences in effort provision directly translate into differences in efficiency.

[7] The comparison of first period effort choices is based on individual observations. Unless otherwise noted, all other tests use matching group averages as independent observations. Reported p-values are always two-sided.

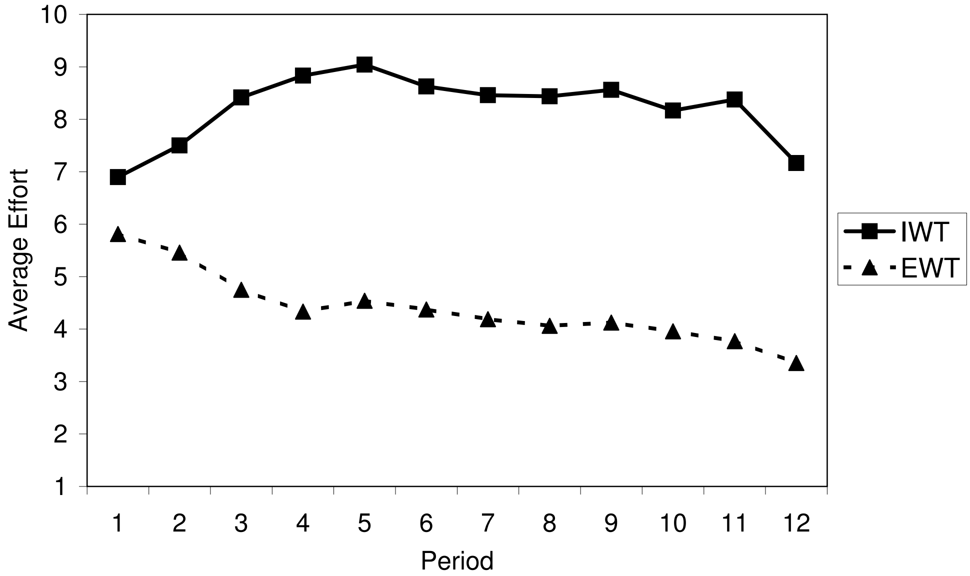

Figure 2.1: *Average effort per period. The effort is aggregated per period over all matching groups.*

> **Result 1:** *The two payment modes exhibit strong differences with respect to the performance they elicit: agents who are paid equal wages exert significantly lower efforts than agents who are paid individually. This results in a much higher efficiency under individual wages.*

Both, the agents and the principals benefit from the increase in efficiency. The average profit per period of a principal is 56 in the EWT compared to 100 in the IWT (Mann-Whitney test: $p < 0.01$), while an agent on average earns 10 under equal wages vs. 17 under individual wages (Mann-Whitney test: $p < 0.01$).

2.3.2 Wage Setting and Monetary Incentives

The strong difference in effort choices suggests that the degree to which gift exchange can mitigate moral-hazard problems depends on the payment mode that is used. Wage equality hampers efficiency, and we hypothesized above that this might be due to horizontal fairness concerns. However, performance differences might also be driven by differing monetary incentives across treatments. To rule this out, we now take a closer look at principals' wage setting and the resulting monetary incentives for the agents.

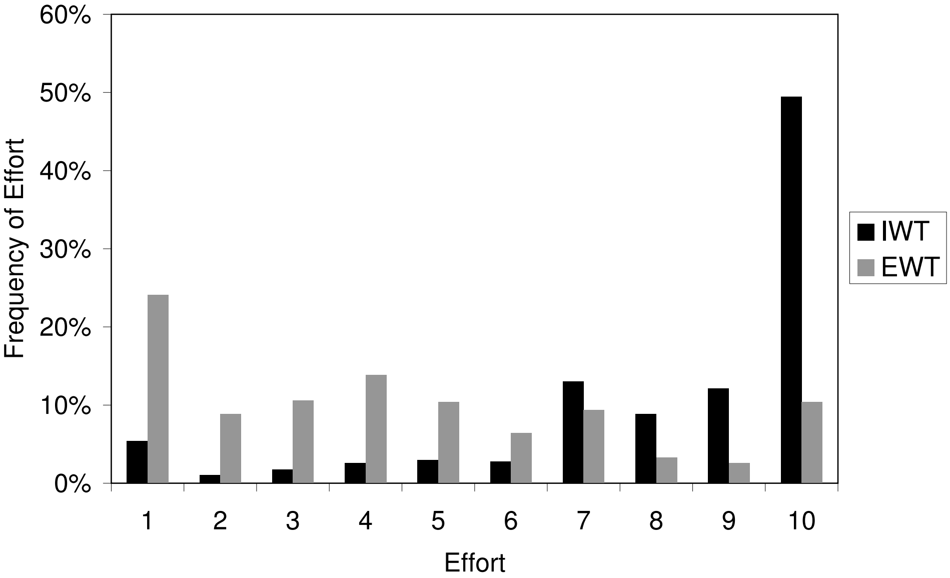

Figure 2.2: *Frequency of effort choices.*

Figure 2.3 plots the average wage per effort level in the two treatments. For both treatments we take the wage paid by the principal for each individual effort decision and calculate averages for a given effort level. The graph exhibits the upward sloping effort-wage relation of many gift-exchange experiments. For example, an agent in the equal wage treatment who exerts an effort of 1 receives on average a wage of 6.3 while an agent exerting an effort of 10 receives an average wage of 30.3. In the individual wage treatment, the corresponding wages are 1.7 and 39.5.[8] The effort-wage relation indicates that gift exchange indeed occurs between principals and agents. In both treatments, higher effort levels are reciprocated with higher wages.

[8]Since principals in the EWT have to pay the same wage to both agents, an interesting question concerns how they choose this wage when confronted with a low and a high effort. To answer this question, we assume that the wage-effort relation from IWT reflects the "true" wage-setting preferences of principals because wage choices were not constrained in this treatment. We regress wages on effort in the IWT and calculate predicted wages for all possible levels of effort. We then calculate the differences between actual wages paid in EWT and these predicted wages. This analysis shows that the actual wage in EWT is very close to the average between the predicted wage for the higher and lower effort in IWT (the detailed analysis is available upon request). This means that principals in the EWT weight the higher and lower effort about equally when deciding on the wage payment.

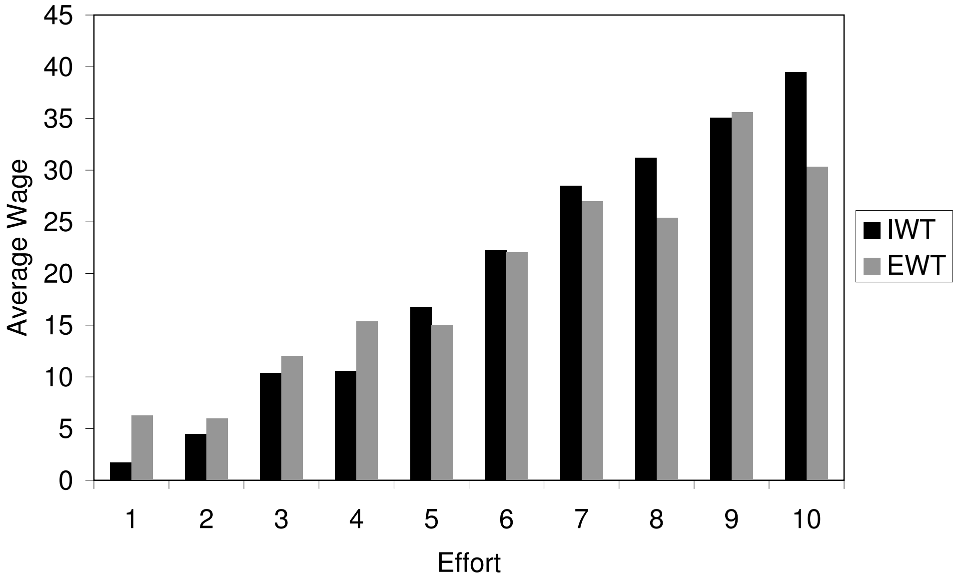

Figure 2.3: *Average wage for a given effort.*

> **Result 2:** *Principals reward higher effort levels with higher wages in both treatments.*

Reciprocal behavior of the principals generates monetary incentives for the agents. In order to calculate the monetary incentives entailed in principals' wage decisions, one has to take into account agents' cost of effort exertion (see Table 2.1). Qualitatively, this does not change the picture of the effort-wage relation: higher effort levels seem to lead not only to higher wages, but also to higher profits for the agents. To check this in more detail, we estimate an OLS-model where we regress the agent's profit per period π_{A_i} on his effort level e_i and a constant. To account for potential differences between treatments we include a treatment dummy IWT, and an interaction term of the treatment dummy and the agent's effort. IWT equals 1 for the individual wage treatment and 0 for the equal wage treatment. Reported robust standard errors are adjusted for clustering within matching groups. The estimation results are shown in Column 1 of Table 2.3. The coefficients indicate that the effort-profit relation is indeed positive in both treatments. On average, an additional unit of effort increases the agent's profit under equal wages by 1.031 points. This coefficient is weakly significant. In the individual wage treatment the effort-profit relation is slightly steeper: an effort increase of 1 leads to an increase in agent's profit of 1.804 points (1.031 + 0.773). The difference between treatments, however, is not significant.

Dep. Variable	π_{A_i}	π_{A_i}
e_i	1.031*	0.854**
	(0.535)	(0.348)
$IWT \times e_i$	0.773	0.995*
	(0.615)	(0.469)
cons	5.927**	-5.815***
	(2.614)	(1.523)
IWT	-3.744	11.004***
	(3.235)	(3.274)
e_j		2.774***
		(0.280)
$IWT \times e_j$		-3.178***
		(0.403)
N. Obs.	576	576
R^2	0.100	0.238

Table 2.3: *Profit regressions. Robust standard errors are adjusted for clustered matching groups and are given in parentheses. For each firm, one observation per period is included in the analysis. The dummy "IWT" is equal to 1 for the individual wage treatment. Significance at the 10%, 5% and 1% level is denoted by *, **, and ***, respectively.*

We also estimate a second model where we control for the co-worker's effort e_j (see Column 2 of Table 2.3). The results indicate that the co-worker's effort choice has a substantial influence on an agent's profit under wage equality while it has a negligible influence if individual wages are paid. An increase in agent j's effort increases agent i's profit in a given period by 2.774 points in the EWT, while the (insignificant) influence in the IWT is -0.404 ($= 2.774 - 3.178$). However, it is still individually profitable for the agents to exert high efforts in the EWT. An additional unit of (own) effort increases the agent's profit by 0.854 points.[9] Our

[9]One could object that subjects in the experiment did not have access to the analyses we just presented, because these are "ex-post" examinations while subjects only observed behavior and

findings concerning agents' monetary incentives can thus be summarized as follows.

> **Result 3:** *The wages paid by principals imply similar monetary incentives in both treatments. A higher effort level leads to a higher profit in both treatments.*

2.3.3 The Importance of Equity

In light of the previous result, the strong differences in actual efforts and especially the low effort levels under equal wages are remarkable and stress the significance of non-pecuniary motivations for agents' performance. Agents under equal wages predominantly choose low efforts, thereby foregoing considerable profits. Apparently, equal wages are not reconcilable with agents' horizontal fairness considerations. On the other hand, agents under individual wages provide very high effort levels. Thus, *aggregated* behavior is consistent with the predictions of equity-concerned agents. We therefore focus our analysis of non-monetary motivations on the question whether *individual* behavior is in line with a concern for the fulfillment of the norm of equity.

Agents' Reactions to Norm Violations

We first analyze how agents react to a violation of the norm of equity. Equity theory argues that agents experience distress from inequity and take action to reduce it—which in our setup means to increase or decrease the individual effort. The direction of the effort adjustment should depend on the type of norm violation. An equity-concerned agent who works more but does not receive a higher payoff than his co-worker faces a *disadvantageous norm violation*. To restore equity, he can only

---

outcomes of their previous groups. We therefore calculate the profit-maximizing effort level for each agent in each period based on the information this subject actually has. If we assume that agents choose the effort level that was on average the most profitable of all effort levels they have observed so far, the calculations show that agents in the EWT could have increased their efforts and profits considerably even by using only their limited information. In the last period, the average profit-maximizing effort level exceeds the average actual level in that period by 61%. By contrast, the average actual effort levels of subjects in the IWT are very close to the profit-maximizing levels.

	Effort Down	Effort Constant	Effort Up	N. Obs.
EWT				
No Violation	19.1 %	54.4 %	26.5 %	68
Adv. Violation	12.2 %	43.5 %	44.3 %	230
Disadv. Violation	52.6 %	33.9 %	13.5 %	230
Total	30.7 %	40.7 %	28.6 %	528
IWT				
No Violation	19.2 %	51.8 %	29.0 %	448
Adv. Violation	45.0 %	27.5 %	27.5 %	40
Disadv. Violation	35.0 %	57.5 %	7.5 %	40
Total	22.3 %	50.4 %	27.3 %	528

Table 2.4: *Frequency of effort reactions.*

decrease his effort. Analogously, his co-worker who exerts a lower effort and earns a higher profit faces an *advantageous norm violation* and should increase his effort.[10]

Table 2.4 shows how often agents decrease, increase or do not change their effort from period t to $t+1$ after they experienced no, an advantageous or a disadvantageous norm violation in period t. The top panel of Table 2.4 reports data for the equal wage treatment. When the norm is fulfilled, most agents keep their effort constant (54%) and slightly more agents increase their effort than decrease it. After experiencing an advantageous norm violation, agents tend to increase their effort (44%) and only few reduce it (12%). The opposite is true after a disadvantageous norm violation: the majority of agents decrease their effort (53%) and only few increase their effort in the following period (14%). In line with equity theory these numbers suggest that agents change their effort provision in the direction that makes a violation less likely to occur in the next period.

Behavior in the individual wage treatment (bottom panel) is very similar to behavior in the EWT for the cases of no violation and disadvantageous violations.

[10] More precisely, an advantageous norm violation comprises all cases when efforts are equal but profit is higher, or when effort is lower but profit is not. A disadvantageous norm violation occurs if efforts are equal but profit is lower, or if effort is higher but profit is not.

When the norm is not violated agents mostly keep their effort unchanged. After a disadvantageous norm violation efforts are decreased rather than increased, as in the EWT. The only difference between the treatments is observed when agents experience an advantageous norm violation: agents in the IWT tend to decrease their effort while the EWT agents tend to increase it in this case.[11]

The pattern of individual reactions to norm violations indicates that agents care about equity; we therefore check next how often norm violations occur in the two treatments. We expected to see more norm violations in the EWT than in the IWT, because the equal wage institution forces principals to set wages that are not in line with the norm of equity whenever agents exert different efforts. This is indeed what we observe. While the norm is violated in 87% of all cases (460 out of 528) in the EWT, the figure for the IWT is only 15% of all cases (80 out of 528). Thus, even if individual reactions in a given situation are similar, agents in the EWT are far more often exposed to norm violations than agents in the IWT. Principals in the IWT seem to understand quite well that agents care about equity and use the possibility to set different wages in a sophisticated way. If efforts differ, they reward the more hard-working agent with a higher wage in 90% of these cases. If agents exert the same effort, principals pay equal wages in 90% of the cases.

> **Result 4:** *Agents mostly react to disadvantageous violations of the norm of equity by reducing their effort and by increasing it after an advantageous norm violation. The norm of equity is far more often violated in the equal wage treatment.*

So far we have seen that agents' reactions are largely in line with the hypotheses of equity theory and that treatments differ with respect to the frequency of equity-norm violations. Yet, this is not sufficient to explain the treatment effect, since a

[11]We checked the robustness of the reaction patterns in several ways. For example, it could be that agents react differently to norm violations if they are paid very high or low absolute wages. However, performing the analysis only for agents receiving a wage out of the top or bottom quartile of the ex-post wage distribution does not alter the result. An implicit assumption of our analysis is that the gift-exchange relation is generally intact between principal and agent, i.e., that agents exert a non-minimal effort and that principals pay a positive wage. The results do not change if one restricts the analysis to these cases. Also if one defines gift exchange as requiring the agent's profit to be positive, i.e. $w_i > c(e_i)$ instead of $w_i > 0$, the results are very similar.

norm violation is always advantageous for one agent and at the same time disadvantageous for the other one. If both agents adjust their effort in a similar way but in opposite directions the adjustments will cancel out. However, previous evidence suggests that reactions to a disadvantageous norm violation are stronger than reactions to an advantageous one (e.g., Loewenstein et al. 1989, Mowday 1991, Thöni and Gächter 2008). If this is the case, norm violations could explain the downward trend in the EWT and the treatment difference in effort provision.

Figure 2.4 shows the average magnitude of changes in effort provision from period t to period $t+1$ after an agent experienced no norm violation, a disadvantageous or an advantageous norm violation in period t. The width of the bars corresponds to the number of observations in the respective category (cf. last column of Table 2.4). When the equity-norm is not violated agents tend to keep their effort constant or even slightly increase it. After a disadvantageous norm violation, agents in the EWT react strongly. They decrease their effort by 1.30. Their co-worker, experiencing an advantageous norm violation, increases his effort but not as strong. He raises his effort by only 0.75. The difference is statistically significant (Wilcoxon test of the absolute values: $p = 0.01$). In the IWT, both groups of agents experiencing a norm violation decrease their effort. The strength of reactions indicate that agents suffer more from a disadvantageous norm violation than from an advantageous one. This results in an overall decrease of efforts after a norm violation.

> **Result 5:** *Agents' reactions to a violation of the norm of equity are asymmetric: the negative reaction of the disadvantaged agents is stronger than the reaction of the advantaged agents. This asymmetry in agents' reactions results in an overall negative time trend in efforts for the EWT and in the strong treatment difference in effort.*

The analyses above suggests that agents care about equity and experience the equal wage scheme as unfair. Interestingly, even the principals consider the equal wage scheme as less fair. In the post-experimental questionnaire, principals were presented three hypothetical game situations that included effort choices, wage choices, and the resulting payoffs for all players. They were asked whether they considered the resulting allocation as just. One of the three situations reflected their own aver-

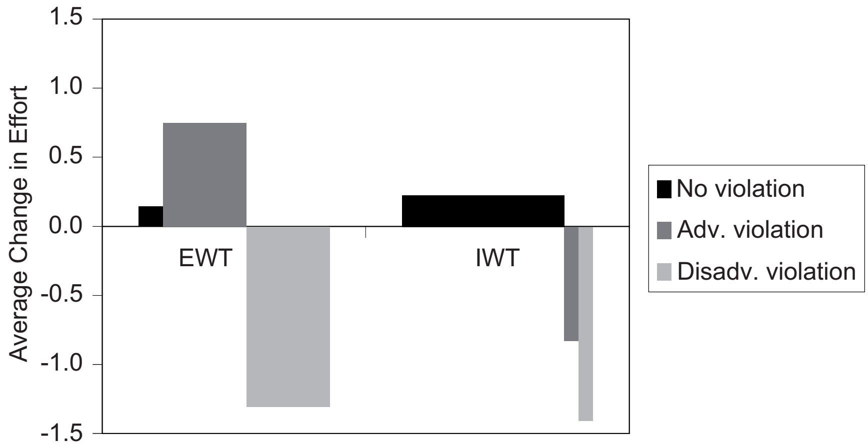

Figure 2.4: *Magnitude of effort reactions. The average change in effort from period t to period $t+1$ is shown given that the agent experienced no norm violation, an advantageous violation or a disadvantageous norm violation in period t. The width of the bars corresponds to the number of observations.*

age behavior in the experiment.[12] The principals did not know that they were facing their own past decisions when answering the question. 63% of the principals in the IWT considered their own decisions fair while only 38% of the principals in the EWT shared this view (Mann-Whitney test on matching group shares: $p = 0.03$).

Simulation with Equity-Concerned Agents

We demonstrated above that horizontal fairness concerns shape agents' behavior under the two payment schemes. In combination with the frequent violations of the norm of equity in the EWT we are able to explain the performance differences across treatments. In order to further illustrate how institutions and equity-concerns interact, we take our previous findings on agents' period-to-period reactions and link them to the aggregate dynamics in the experiment. We do so with a simulation in which all agents are assumed to derive utility from money, but to also suffer whenever the equity principle is not met. When deciding about their effort in a given period,

[12] This situation was constructed as follows: We calculated the average effort of the higher-effort and of the lower-effort providers that the principals actually faced during the experiment. We then took the average of the wages the principals paid to the two groups. Finally, we calculated hypothetical payoffs for all three "average" players by considering the costs of the average efforts.

the simulated agents compare their effort and profit in the previous period with the effort and profit of their co-worker in that period. According to the comparison along these two dimensions, four reactions can be distinguished for the simulated agents. (i) For an agent who had a higher effort and a higher profit, the norm of equity is fulfilled and the pecuniary comparison is also advantageous for him, so he keeps his effort constant. (ii) For an agent who exerted a lower effort and got a lower profit, the norm is satisfied but profit maximization is not, thus he partly adjusts his effort in the direction of his co-worker's effort, i.e., he chooses an effort $(e_{i,t} + e_{j,t})/2$. (iii) An agent with higher effort and lower profit feels distressed as he suffers from a disadvantageous norm violation. He adjusts his effort fully and chooses $e_{i,t+1} = e_{j,t}$. (iv) Finally, for an agent with lower effort and higher profit the norm violation is advantageous, thus the resulting utility is higher than in case (iii). He chooses an effort $(e_{i,t} + e_{j,t})/2$. The reactions in cases (i) to (iv) are in line with the period-to-period reactions presented in Table 2.4 and Figure 2.4.

In the simulation, we use actual effort data from the experiment only for the first period. The subsequent effort decisions are based on the simulated profits and simulated efforts of the previous period. The simulated principals pay the average wage for a given effort (IWT) or the average wage sum for a given effort sum (EWT) as calculated from the experimental data. Profits are then calculated as wage minus cost of effort exertion. We use the same matching protocol as in the experiment.

Figure 2.5 shows how effort choices evolve over time in the experimental data and in the simulations. The simulations *'EWT sim'* and *'IWT sim'* trace the real data very well and are able to reproduce the large effort difference between treatments. In the individual wage simulation, efforts increase like the real efforts although the slight downward trend in the second half of the experiment cannot be reproduced. Efforts in the equal wage simulation constantly decrease down to an effort level slightly above 3 in the final period. This pattern is very similar to the dynamics in the real data.

Note that the pivotal agent is different between the simulated treatments: in the equal wage simulation the norm of equity is violated when agents choose different effort levels. In these cases, the agent with the higher effort will fully adjust his

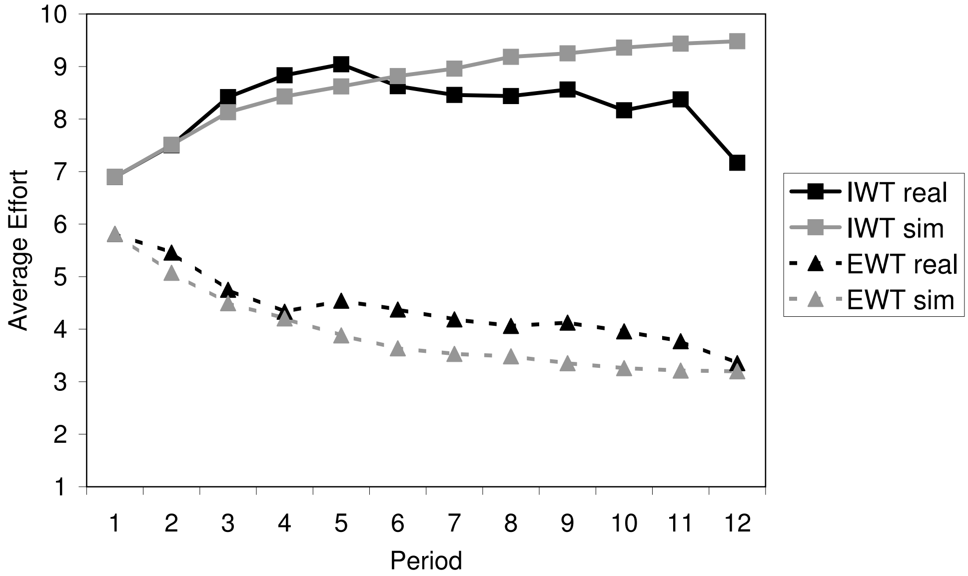

Figure 2.5: *Simulated efforts of agents adopting to equity-norm violations.*

effort in the direction of his co-worker's effort while the co-worker will increase his effort level only to the average effort of the last period. In the EWT simulation, the average effort therefore converges to the lowest first period effort as agents are subsequently re-matched: the low-effort providers are pivotal. By contrast, in the IWT the high-effort providers have the decisive impact on the overall outcome. The norm of equity is mostly fulfilled in the IWT. Thus, the agent with the higher effort keeps his effort constant while his co-worker adjusts his effort. The average effort therefore converges to the highest first period effort. We will analyze this point in more detail in the next section.

> **Result 6:** *Simple simulations based on agents who have preferences for money and equitable treatment are in line with the efforts observed in the experiment and are able to reproduce the observed treatment effect.*

2.3.4 Dynamics of High-Effort and Low-Effort Providers

As already seen in Figure 2.2, subjects exhibit a substantial degree of heterogeneity with respect to effort provision. In the following, we analyze if the agents who are most or least willing to exert effort are affected differently by the two payment modes at hand. A common informal argument claims that equal wages will be especially detrimental to the motivation of high performers but clean empirical evidence is scarce. Furthermore, it is unclear how weakly motivated agents react to equal or individual wages. We also address the question whether high and low performers impact the overall results differently in the two treatments. The simulations presented in the previous section suggest that this could indeed be the case: in the EWT simulation, the low-effort providers are decisive for the final outcome while it is the high-effort providers in the IWT simulation.

To analyze these questions in the experimental data we classify agents according to their effort decision in the *first period*. We define the agent with the highest first-period effort in each matching group as "high-effort provider" and the agent with the lowest effort as "low-effort provider". This type definition is chosen because when agents decide on their effort in the first period, they do not have any information about the behavior of other subjects and all learning and coordination processes occur after this initial effort choice. Thus first-period effort is likely to be a good proxy for the intrinsic willingness of a specific agent to exert effort. If some of the subjects are intrinsically inclined to exert high efforts they should show up in the group of high-effort providers. In contrast, if some of the subjects are intrinsically inclined to exert low efforts they should show up in the group of low-effort providers.

In Figure 2.6 we follow the high-effort providers and low-effort providers in both treatments and show their effort decisions over time. In the first period, the groups of high-effort providers and the groups of low-effort providers are close together across treatments.[13] This changes completely over the course of the 12 periods. In the individual wage treatment, high-effort providers continue to provide high effort levels. Low-effort providers increase their efforts dramatically up to the level of the high-effort providers and even higher in the last periods. In the equal wage treatment, the dynamics are reversed. Here, the low-effort providers keep their

[13]In the first period, effort levels are not significantly different between treatments for high-effort providers (Mann-Whitney test: $p = 0.14$) while they are close together but different for the low-effort providers (Mann-Whitney test: $p = 0.03$). Within treatments, the high-effort and low-effort providers choose statistically different effort levels in the first period (Wilcoxon signed rank test: $p = 0.01$ (IWT), $p = 0.01$ (EWT)).

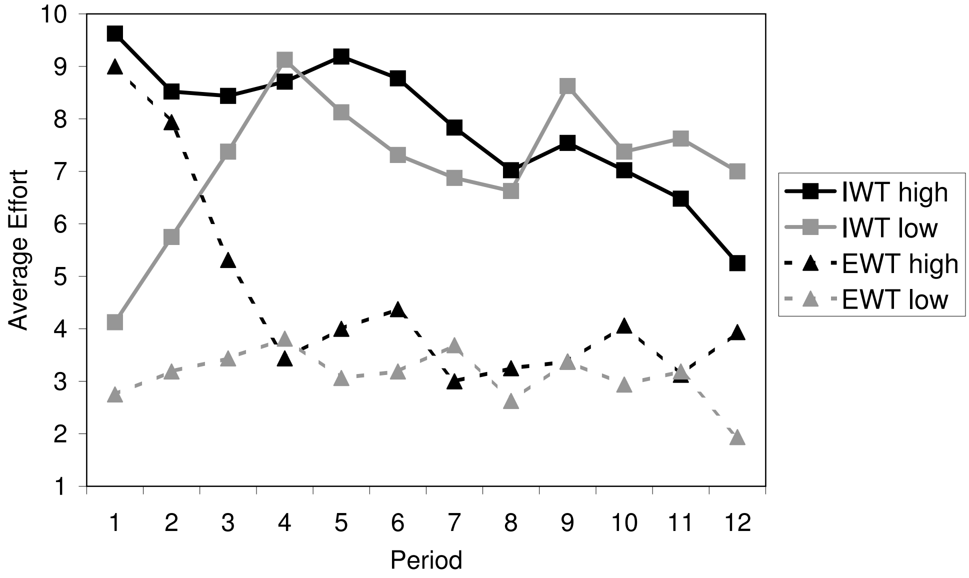

Figure 2.6: *Effort decisions of high-effort and low-effort providers. In each matching group, the agent with the highest (lowest) effort in the first period is defined as the high (low)-effort provider.*

effort provision constant and the high-effort providers reduce their efforts to the level of the low-effort providers. In the last six periods, effort levels are not different within treatments (Wilcoxon signed rank test: $p = 0.67$ (IWT), $p = 0.78$ (EWT)) while they differ between treatments (Mann-Whitney test: $p < 0.01$ (high-effort providers), $p < 0.01$ (low-effort providers)). Put simply, the "good" agents push the "bad" agents up under individual wages while under equal wages the "bad" ones pull the "good" ones down.

These dynamics underline the importance of the different non-monetary motives induced by the two wage setting institutions. Remember that agents face similar monetary incentives in both treatments, but wage equality often violates the norm of equity. Agents in this treatment who are in principle willing to exert high levels of effort get frustrated and lower their efforts. On the contrary, under individual wages where the norm of equity is intact, good performance spreads. These results suggest that choosing a wage scheme also influences the social dynamics between the agents. In our experiment, individual wages lead to positive dynamics since agents orientate themselves by the most hard-working agents. In contrast, the equal wage scheme focuses agents' attention on the least motivated agents.

Result 7: *The pivotal agent is different between treatments: in the IWT, agents who initially provide low effort align with the high-effort providers over time. In the EWT, agents who initially provide high effort align with the low-effort providers over time.*

2.3.5 The Role of Intentions

So far, we interpret our results as supporting the notion that subjects care about the norm of equity. However, by design our treatments necessarily differ in the number of instruments that a principal has at hand. In the EWT, principals only choose a single wage—whereas principals in the IWT decide on two wages and consequently can tailor reactions individually to agents' preceding choices. Therefore, agents might attribute a different degree of intentionality to principals' decisions: In the EWT, the role of intentions is limited to the level of the wage. The IWT contains an additional element of intentionality because principals also decide on relative wages and consequently whether the equity norm is fulfilled or violated. In light of the literature that stresses the behavioral importance of intentions in situations of reciprocal interaction (e.g., Dufwenberg and Kirchsteiger 2004, Falk et al. 2005), there is thus a potential alternative explanation for our treatment effect. In other words, one might speculate that the difference is not caused by the different frequency of norm fulfillment *per se*, but rather by the additional element of intentionality.[14]

To test this alternative explanation, we conducted an additional control treatment (*wage level treatment* or WLT) that clearly isolates the effect of norm fulfillment on agents' effort choices. As in the EWT, principals in the WLT only choose a single wage. The other agents' wage is then exogenously set by a computer program such that the equity norm is always fulfilled, i.e., agents who exerted a higher effort than their co-worker automatically receive a higher payoff. Importantly, this implies that the fulfillment of the equity norm is not attributable to principals' decisions. Except for this change of the wage-setting institution, the instructions and the experimental design were identical to the previous treatments. The 72 subjects who participated in the four additional sessions had not previously taken part in the

[14] We thank Patrick Bolton and an anonymous referee for pointing this out.

IWT or the EWT.

The specific equity norm implemented in the WLT experiments dictates proportionality between agents' monetary payoffs and efforts. We chose this "equity formula" as it is probably the most prominent formulation of the equity principle (see Section 2.2.2). Given a principal's decision for the low-effort agent, the wage for the high-performing agent is exogenously fixed such that both agents receive the same payoff per unit of effort provided, i.e., $\frac{\pi_{\text{low}}}{e_{\text{low}}} = \frac{\pi_{\text{high}}}{e_{\text{high}}}$ holds. For example, if the principal observes efforts of 2 and 6 and sets the wage for the low-effort provider to be 5, the payoff of this agent is $5 - c(2) = 4$ (compare Table 1). Following the equity formula, the payoff of the high-effort provider will then automatically be set to $\frac{4}{2} \cdot 6 = 12$; which implies a wage of 20 after taking the cost of providing 6 units of effort into account.[15]

The wage-setting institution in the WLT is not meant to be an analog of institutions found in actual labor markets, as it is the case for the IWT and the EWT. It exogenously implements the incentive structure that is endogenously created by principals in the IWT.[16] If we observe similar efforts in the WLT as in the IWT we can rule out intentions as an explanation for the difference between our two main treatments, IWT and EWT.

Figure 2.7 compares agents' mean effort choices over time for all three treatments. As can be seen, the exogenous implementation of the equity norm suffices to elicit high efforts from the agents. The average effort difference between the WLT and the

[15]The equity formula leads to counterintuitive implications whenever negative values for the inputs or outcomes are possible. Therefore, if in our experiment the efforts differ and the principal's choice of w_{low} implies $\pi_{\text{low}} \leq 0$, the other agent's wage is instead set such that $\pi_{\text{high}} = \pi_{\text{low}} + 5$. This guarantees that a norm of equity is fulfilled for all possible wage-effort combinations. Nevertheless, the high-effort agent still faces the risk of making losses whenever the low-effort agent gets a negative payoff.

[16]As shown in Result 3, the monetary incentives in the IWT imply that profit-maximizing agents should provide non-minimal effort levels. As a consequence of exogenously implementing these implicit incentives in the WLT, new subgame-perfect Nash equilibria necessarily arise. Our focus of interest, however, rests on the comparison of the observed behavior across treatments rather than on comparing behavior to the game-theoretical equilibrium predictions. For a similar approach of "exogenizing" endogenous incentives to test for the impact of intentions, compare for example Blount (1995), Charness (2004), or Cox (2004).

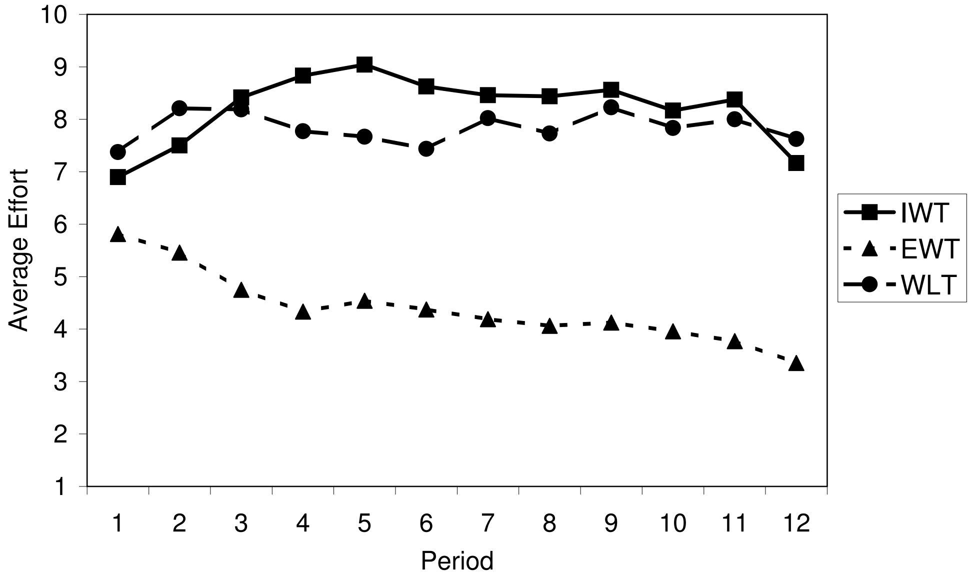

Figure 2.7: *Average effort per treatment.*

IWT of only 0.36 is insignificant (Mann-Whitney test: $p = 0.83$). Compared to the EWT, efforts are on average 3.44 units higher in the wage-level treatment (Mann-Whitney test: $p < 0.01$). As in the IWT, efforts do not decrease over time in the WLT (Wilcoxon test for periods 1–6 against 7–12: $p = 0.44$). Also the distribution of efforts in the WLT closely resembles the one in the IWT. Under both treatments, the modal choice is the provision of maximum effort. In the WLT, an effort level of 10 is chosen in 46.5% of all cases, compared to 49% in the IWT.

> **Result 8:** *The wage level treatment shows that the treatment difference between the IWT and the EWT is not caused by the fact that principals can set two wages instead of one* per se. *Differences in equity norm fulfillment—independent of intentionality—are the driving force behind agents' performance.*

Taken together, the results from the additional control treatment corroborate our previous findings. They suggest that the observed performance differences are not driven by the differing degree of intentionality across treatments. This, of course, does not imply that intentions are unimportant in general; however in our setup, treatment differences are almost exclusively caused by equity considerations.

2.4 Concluding Remarks

In this chapter, we studied the interaction of gift exchange with different payment modes; more specifically, we analyzed how horizontal fairness concerns of employees affect the effectiveness of gift exchange as a contract enforcement device. In our experiment, one principal is matched with two agents. The principal pays equal wages in one treatment and can set individual wages in the other. The use of equal wages elicits substantially lower efforts and efficiency in spite of similar monetary incentives: exerting high effort pays off under both wage schemes. The strong treatment difference is driven by subjects' preferences for horizontal equity and the fact that the equity principle is frequently violated in the equal wage treatment. This is not the case in the individual wage treatment, as principals set wages almost always in line with the norm of equity. The results of a control treatment show that indeed norm fulfillment *per se* and not different degrees of intentionality are the driving force behind agents' behavior.

Our results have a number of implications, both for the advancement of existing theories and for the design of wage schemes in practice. First of all, while it is well-known that equal wages can distort monetary incentives, in our experiment they are efficiency decreasing *even though* individuals' monetary incentives are qualitatively not affected. Rather, equal wages oftentimes lead to situations which are considered as unfair by the workforce. This holds in particular because agents are heterogeneous and equal wages violate the equity principle whenever workers differ in their performance. It may thus be oversimplifying to argue that equal wages lead to less envy and therefore higher work morale, as it is frequently done in the political discussion.

In this regard, it is doubtful whether strict wage equality can be reconciled with the use of reciprocal gift exchange to enforce incomplete contracts. Our findings suggest that adherence to the norm of equity is a necessary prerequisite for a successful gift-exchange relation. Consequently, the wage setting institution must provide principals with means to account for possible differences in agents' behavior, e.g., to individually reward agents who outperform their co-workers. The performance of agents in the individual wage treatment and in the wage level treatment shows how effective gift exchange can be, as long as horizontal equity concerns are respected: although explicit contract enforcement is absent, 80% of the possible efficiency gains are realized.

In practice, the discretion to fulfill the norm of equity does not have to be in monetary terms. Perks and non-monetary benefits like extra vacation or awards can be useful devices to motivate workers in this context. These instruments become especially important when it is not possible to wage discriminate on a given hierarchical level, e.g., because the firm's internal pay structure, agreements with a union or legistlation dictate wage equality.

The results in this chapter should not be interpreted as arguments against wage equality in general. They rather suggest that equal wages come at a cost that has to be weighed against their potential benefits. For example, equal wages are easier to implement than individual wages, and they may encourage peer monitoring and collaboration. The relative importance of these costs and benefits (and also the impact of the workforce's social preferences more generally) is likely to depend on the details of the institutional setting. These include the production technology, the information structure, and the organizational design of the firm. In this chapter we presented results for one such setting. Our design provides a simple and parsimonious framework that can successively be enriched to study these aspects in future research.

Chapter 3

Contract Enforcement and Involuntary Unemployment

3.1 Introduction

Many labor relationships are characterized by contractual incompleteness, in particular since employment contracts often specify employees' obligations only imprecisely. Probably the most important reason for leaving contracts incomplete is the difficulty to monitor and verify employee performance or work effort. When this is the case, firms cannot stipulate fully contingent contracts but have to rely on other, implicit means to elicit work effort. Several such implicit contract enforcement mechanisms have been discussed in the literature, ranging from voluntary bonus payments to self-enforcing, relational contracts based on an implicit threat to dismiss workers who have been caught shirking (e.g., MacLeod and Malcomson 1989, Baker et al. 1992, MacLeod and Malcomson 1998). Other papers have argued that social preferences could help to overcome moral hazard problems. Workers who feel treated fairly, e.g., because their employer generously shares the rents from production, might voluntarily provide high work effort in return (Akerlof 1982, Fehr et al. 1993).

While implicit contracts can provide strong performance incentives, they can have detrimental consequences for the labor market as a whole. In particular, implicit contracting can give rise to involuntary unemployment. For instance, when fairness concerns are important, it can be optimal for firms to dismiss workers rather than cut wages in times of economic downturns (Bewley 1999). More generally, firms might ration jobs because the requirement to pay fair wages could render less productive jobs unprofitable. Finally, the presence of unemployed workers could itself be a prerequisite to elicit effort of those employed. Higher unemployment levels increase the cost of job losses since, c.p., the job acquisition rate for unemployed workers decreases. Thus, unemployment might be necessary to provide sufficient performance incentives for employed workers who want to avoid losing their jobs (Shapiro and Stiglitz 1984, MacLeod and Malcomson 1989).

In this chapter, we empirically analyze the relationship between contract enforcement and the emergence of unemployment. We address the following questions: Does the absence of third party contract enforcement have a direct impact on the level of involuntary unemployment? Is this influence due to the implicit contract enforcement instruments used by firms? In particular, do firms ration jobs when effort is not verifiable although full employment would be (technologically) efficient? Do they engage in repeated, long-term employment relationships and do they pay "fair" wages? How does unemployment influence workers' effort provision?

We study these questions in an experimental labor market where we exogenously vary the verifiability of work effort. In the market, firms and workers interact during multiple market periods. All firms share the same production technology, which exhibits decreasing returns to scale from labor, but ensures that full employment is technologically efficient. In each market period, firms can hire up to two workers whereas workers can accept exactly one contract offer. We concentrate on simple, but frequently observed employment contracts which stipulate a fixed, non-contingent wage payment and a desired level of work effort. In order to study rent-sharing as well as the determinants of dismissal or re-employment in relational contracts, we allow firms and workers to interact repeatedly and build up long-term work relationships.

Our two treatment conditions only differ in the degree to which work effort is verifiable and, hence, explicitly enforceable. In a control treatment, concluded contracts are explicitly enforced, i.e., a worker's effort has to be equal to the contractually agreed upon effort level (*C treatment*). By contrast, in our main treatment (*IC treatment*) effort is observable to firms and workers, but not verifiable. Firms therefore have to rely on implicit contracts. Backward induction together with the assumptions of standard neoclassical labor market models predict that treatments will differ in terms of workers' performance (i.e., first best work effort will be elicited in the C treatment, whereas effort will be minimal in the IC treatment), but that maximum employment will be achieved in both treatments.

However, when workers' performance depends on whether they feel treated fairly by their employer (i.e., when workers have social preferences), firms in the IC treatment might find it profitable to ration jobs, rely on repeated interaction with specific workers, and pay strictly positive rents to those workers. This could lead to market outcomes where firms succeed in eliciting above minimum effort levels also in the IC treatment. However, this efficiency increase in terms of effort might come at the cost of involuntary unemployment. High wage payments and job rationing in the IC treatment could lead to higher unemployment as compared to the treatment with explicit contract enforcement.

Our experiment yields the following findings. First, unemployment is indeed much higher when third party contract enforcement is not feasible. More importantly, unemployment in this treatment is involuntary, being caused by the firms' employment and contracting policy. Most firms indeed pay high wages but offer fewer vacancies than possible. By contrast, wages are close to the market clearing level when effort is explicitly enforced. Firms do not ration jobs, and unemployment is very low. Moreover, unemployment in this treatment is mostly voluntary, being caused by workers who do not accept existing contract offers.

The firms' employment policy in the IC treatment, however, succeeds in eliciting high efforts from the employed workers: more than 50% of workers choose effort levels close to the maximum and only 12% of workers shirk and provide minimum work effort. We observe that many firms employ a specific worker over several periods when effort is not explicitly enforced. These long-term work relationships are characterized by high wages and high effort levels. Providing high efforts in response to high wages is profitable for the worker since firms do not rehire workers who were caught shirking in previous periods. At the same time, being unemployed entails considerable costs for workers. Due to firms' job rationing, the prevalence of long-term work relationship, and the high level of unemployment, the job acquisition rate for unemployed workers is very low.

The IC treatment also allows us to deeper analyze the determinants of firms' success in an environment where contracts are not explicitly enforceable. Interestingly, profits are initially higher for firms who ration jobs, but do not depend on firm size in later periods of the IC treatment. As more and more firms decide to ration jobs, unemployment rises, and therefore the pressure to perform increases for employed workers. Ultimately, work effort (and therefore firm profits) also increases in firms that employ more than one worker. However, we find that employment relationships in these firms differ strongly from the ones in successful one-worker firms. While one-worker firms are characterized by long-term employment relationships, high wages and high worker rents, the successful two-worker firms offer low wages and positive, but lower rents to the workers. Moreover, fluctuation in two-worker firms is much higher. The result resembles an "endogenously segmented" labor market with two forms of employment relationships which are equally successful from the firms' perspective but differ strongly for workers.

Efficiency wage theories have long postulated that the absence of explicit contract enforcement could lead to involuntary unemployment (Shapiro and Stiglitz 1984, MacLeod and Malcomson 1989, Akerlof and Yellen 1990). To the best of our knowledge, we are the first who empirically show a direct, causal link between contract enforcement and the emergence of unemployment. So far, there has been only indirect evidence which suggested that this link might exist. For instance, it has been shown in surveys as well as in economic experiments that fairness concerns in environments characterized by contractual incompleteness can be the cause of wage rigidities (Campbell III and Kamlani 1997, Bewley 1999, Fehr and Falk 1999). In a setup similar to ours, Brown et al. (2004) have shown that the absence of third party contract enforcement leads to a "bilateralization" of trade, where interaction between firms and workers is characterized by long-term employment relationships with high wages and high effort. Unemployment, however, was exogenously given in their experiments and firms had no option to ration jobs.

The paper most closely related to this chapter is the work by Fehr et al. (1996) who also study the emergence of unemployment in a market where effort is not perfectly verifiable. However, their experimental design differs from ours in a variety of important aspects. First, they concentrate on one-shot interactions between firms and workers while our focus is on the emergence of unemployment through fairness concerns and relational contracting. Second, as a shortcut for implementing a threat of dismissal, Fehr et al. (1996) introduce an exogenously given, strictly positive probability that workers who shirk will be caught and have to pay a penalty to their firm. In contrast, we can directly observe dismissal and re-employment decisions by firms. Finally, workers' productivity in their experiment was private knowledge to firms and productivity in some firms was so low that these firms could not profitably offer incentive compatible contracts. In our setup, we observe job rationing *although* it is always possible for firms to offer an incentive compatible contract that yields positive profits.

Complementary evidence to our results that also suggests the importance of efficiency wages comes from the literature on non-compensating intra- and inter-industry wage differentials (e.g., Krueger and Summers 1988, Blanchflower et al. 1996, Abowd et al. 1999, Goux and Maurin 1999). The earlier papers in this literature have observed persistent interindustry heterogeneity in wages for workers with identical observable characteristics. This rejects classical competitive theories of wage determination. More recently, Abowd et al. (1999) and Goux and Maurin (1999) have found that unobserved worker heterogeneity explains a considerable fraction of interindustry wage differentials. However, strong interfirm wage differentials within a given industry persist.

The results reported in Goux and Maurin (1999) indicate that wages strongly increase in firm size and firms' capital intensity and, to a lower extent, in productivity or profitability. Our results for the IC treatment show that higher wages are paid by firms who ration jobs and employ only one worker. This is, however, only seemingly contradictory. According to Goux and Maurin (1999), it is likely that larger firms in their sample face more difficulties to monitor employee performance. Therefore, their findings are in line both with efficiency wage theories and with our results. Comparison of wages *across* our treatment conditions strongly suggests that it is indeed the difference in the verifiability of effort and not firm size per se that matters for differences in wages. Holding contract enforcement constant—i.e., comparing wages *within* the IC treatment—it is the more productive and, under decreasing returns to scale, smaller firms that pay higher wages.

Our finding of an endogenous segmentation of the labor market in the IC treatment also provides interesting new insights into the debate on dual labor markets. It has long been acknowledged that dual labor markets can be an implication of contract enforcement problems and efficiency wages. However, in these models, dual labor markets typically arise as a consequence of differences in monitoring costs across firms (Bulow and Summers 1986) or due to differences in adjustment costs to demand fluctuations (Saint-Paul 1996, ch. 4). In our experiment, segmentation occurs in the IC treatment even though all firms face the same technological constraints. Rather, segmentation is the result of market interactions when third party contract enforcement is not feasible. In response to the non-verifiability of effort, many firms ration jobs, pay high wages and employ specific workers over a long time horizon. As a result, unemployment increases. This allows other firms to successfully employ a "secondary-sector" contracting strategy, involving lower wage payments and tighter conditions for contract renewal. High unemployment disproportionally helps these firms, since they are more likely to hire previously unemployed workers. For a given wage, these workers tend to shirk particularly little—unlike in classical shirking models where unemployment disciplines all workers similarly. The "secondary sector" firms are therefore able to pay lower wages to elicit a certain level of work effort.

The remainder of this chapter is organized as follows: the following section describes our experimental setup. Section 3.3 derives behavioral predictions and Section 3.4 presents the empirical results. Section 3.5 concludes.

3.2 Experimental Design and Procedures

To study the impact of contractual incompleteness on unemployment, we implemented an experimental labor market where we exogenously varied the verifiability of work effort. In the market, firms and workers interacted during 18 market periods. Each of the 18 periods consisted of two stages: a market phase where firms offered employment contracts and hired workers, and a work phase where work effort of employed workers was determined. The experimental treatments differed only in the degree to which work effort in the second phase was third party enforceable. In our main treatment, henceforth called *Incomplete Contracts Treatment* (or *IC treatment*), third party contract enforcement was not feasible and workers thus could depart from the contractually agreed upon effort level. By contrast, the effort level stipulated in the employment contract was explicitly enforced in our control treatment, henceforth called *Complete Contracts Treatment* (or *C treatment*).

3.2.1 The Market Phase

Firms were the contract makers in the market phase. When offering a contract, firms stipulated a non-contingent wage payment w and a desired level of effort $\hat{e}$. To study the relevance of long-term employment relations under the different treatment conditions, firms could make two different types of contract offers: public offers which were available to all workers and could also be observed by all other firms, or private contract offers that were only available to one specific worker. The latter type of contract offers allowed firms to rehire certain workers and interact repeatedly with them. To enable the formation of long-term relations, in the beginning of the experiment each worker and each firm received an identification number (ID) which was held constant throughout the whole experimental session. If an employer wanted to (re)hire a specific worker via a private contract offer, she had to specify the ID of the worker in addition to the wage and desired effort level when entering the contract offer. In this case, only the selected worker was informed about the contract offer, and only this worker could accept the offer.

In a certain market period, each employer could hire up to two workers. As long as none of her contract offers had been accepted in a given period, an employer could make as many private and public offers as she wanted. A worker could accept all contract offers available to him, i.e., all public offers that were not yet accepted and private offers that firms had addressed to him. Once a worker accepted a contract offer, the contract between this worker and the respective firm was concluded. After agreeing on one contract, the worker was not allowed to accept further contract offers in this period. Additionally, all other outstanding offers of the respective employer were removed from the list of available contracts in the moment where one of her contract offers was accepted. The employer could then decide to hire a second worker by entering new contract offers. This market feature was implemented to prevent that an employer who wanted to employ only one worker but entered multiple contract offers had two offers accepted before being able to withdraw her remaining contract offers.

The market phase ended when the maximum number of contracts had been concluded or when all firms had indicated that they did not want to make additional contract offers.[1] At the end of the market period, the worker(s) of a given firm were informed about the contracts concluded by their firm, i.e., each worker received a summary of his own contract terms as well as information on whether and under which conditions his firm had employed a second worker. Providing information on the size of the firm in which a worker is employed can be crucial for workers when contracts are not third party enforceable. If some workers base their work effort on the extent to which their firm shares production rents, e.g., because they respond to “fair” wage payments, knowledge of offered rents, productivity, and firm size is of high importance for workers (see Section 3.2.3 and Section 3.3).

[1]We also implemented a maximum trading time of 200 seconds for each market phase. This constraint was, however, only binding in few occasions (mostly in the C treatment). The impact of the time constraint on the level of unemployment and other market outcomes reported below is therefore limited and confined to the control treatment with explicit contract enforcement (see Section 3.4.1).

3.2.2 The Work Phase

After the end of the market phase, the employed workers entered the second stage of a market period—the work phase. In this stage, actual work effort e was determined. Since effort was contractible in the complete contracts treatment, workers who had accepted a contract offer in this treatment had to comply with the contract terms. The desired effort level $\hat{e}$ stipulated in their contract was thus explicitly enforced, i.e., $e = \hat{e}$ was exogenously implemented by the experimenter. By contrast, work effort was observable, but not verifiable in our main treatment. Therefore, a worker in the IC treatment could choose any feasible level of effort in the work phase, i.e., he could also exert less or more effort than stipulated in his employment contract. Workers' effort choices, together with firms' wage payments, determined material payoffs of firms and workers. Before the next period started, a firm and its worker(s) were informed about actual work efforts and the resulting payoffs for the firm and the workers employed by this firm.

3.2.3 Parameters and Procedures

The participants' roles were randomly assigned at the beginning of the experiment and kept constant throughout all market periods. In every market, we had 17 workers and 7 firms. Since firms could employ at most two workers, this implies that three workers were "exogenously" unemployed in each period (see Section 3.3). A worker's material payoff π_W was given by

$$\pi_W = \begin{cases} w - c(e) & \text{if worker accepted a contract } [w, \hat{e}] \\ 0 & \text{if unemployed} \end{cases}$$

A worker who remained unemployed in a given period received a payoff of 0 points. An employed worker received the wage w specified in his contract and had to bear the cost of the work effort he provided, $c(e)$. The set of feasible efforts and wages was given by $e \in \{1, 2, ..., 10\}$ and $w \in \{0, 1, 2, ..., 100\}$. Effort costs $c(e)$ increased convexly in the level of actual work effort (see Table 3.1).

A firm's material payoff depended on the number of workers hired, the wage(s) paid, and the effort exerted by the worker(s). Firms' production technology was

Effort level e	1	2	3	4	5	6	7	8	9	10
Cost of effort $c(e)$	0	1	2	4	6	8	10	12	15	18

Table 3.1: *Schedule of effort costs.*

characterized by decreasing returns to scale. Specifically, each unit of effort by a worker increased production (and the firm's payoff) by 10 points if only one worker was employed by the firm. If two workers were employed, each unit of effort increased the firm's payoff by 7 points. Additionally, firms had to pay the wages specified in their contracts. The material payoff of a firm, π_F, can therefore be summarized as follows:

$$\pi_F = \begin{cases} 10e_1 - w_1 & \text{if one worker employed} \\ 7(e_1 + e_2) - w_1 - w_2 & \text{if two workers employed} \\ 0 & \text{else} \end{cases}$$

e_1 (e_2) denotes the effort provided by the first (second) worker, and w_1 (w_2) is the wage paid to the first (second) worker employed by the firm. Note that this specification of the production technology implies that efficiency is maximized when two workers are employed and maximum effort is exerted: the second worker's marginal productivity per unit of effort is 4 whereas the marginal cost of effort lies between 1 and 3 points. Payoff functions π_F and π_W, workers' cost schedule $c(e)$ as well as the number of firms and workers in the market were common knowledge.

The experiment was carried out between June and November 2007 at the Bonn-EconLab, the laboratory for economic experiments at the University of Bonn. We conducted five sessions each for the IC treatment and the C treatment. A total of 240 subjects, mainly undergraduate university students from all majors, took part in the experiments. Every subject participated only in one treatment condition. At the beginning of an experimental session, participants received detailed information about the rules and structure of the experiment.[2] The experiment started only af-

[2]Instructions can be found in Appendix C.4. In order to prevent that participants' experience with fairness norms and long-term relationships from labor relations outside the laboratory play a prominent role, instructions were framed in a neutral goods-market language.

ter all participants had answered several control questions correctly. In addition, subjects played one trial period of the market phase to ensure that they understood how to use the computer program. Sessions lasted about 110 minutes and subjects earned on average 25.49 Euro (about 35 USD at the time of the experiment), including a showup fee of 8 Euro. The experiments were computerized using the software "z-Tree" (Fischbacher 2007); subjects were recruited using the online recruitment system by Greiner (2003).

3.3 Behavioral Predictions

3.3.1 Money-Maximizing Behavior of all Players

As indicated in Section 3.2, efficiency is maximized when all firms employ two workers and all workers exert full effort. If players are rational and selfish, and if effort is contractible (C treatment), we should expect that profit-maximizing firms implement the first best outcome and reap all gains from trade: in each period, firms will employ two workers, offering contracts which ask for maximum effort ($\hat{e} = e = 10$) and pay the minimum acceptable wage. As unemployed workers receive zero unemployment benefits, this wage is equal to the cost of the implemented effort, i.e., $w = c(10) = 18$ points. Since there are 7 firms and 17 workers in each market, three workers will remain unemployed by design of the experiment. As we are interested in unemployment which is arising *endogenously* through the market interaction between firms and workers, we will refer to the situation where the maximum number of workers is employed and only three workers remain "exogenously" unemployed as *full employment* in what follows.

Which market outcomes should we expect in the IC treatment under the assumptions of standard neoclassical labor market models, i.e., rationality, money-maximizing behavior of all players and common knowledge thereof? All players know that the market operates for exactly 18 periods—our setup thus constitutes a repeated game of finite length which is solvable by backward induction. Since effort is costly, but not verifiable, workers will choose the minimum effort $e = 1$

in the final period of the IC treatment, irrespective of the wage and desired effort level stipulated in their contract. Anticipating this, money-maximizing firms will pay the minimum acceptable wage for an effort of 1 in this period. This amounts to a wage payment $w = c(1) = 0$ since, again, workers' outside option is to remain unemployed and receive zero payoff.[3] Through backward induction, the outcome of minimal effort (and minimal wage) will also hold in all pre-final periods of the IC treatment.

In sum, predictions differ strongly in terms of contracts offered and effort implemented in the two different market environments. Lower efforts in the IC treatment directly translate into lower market efficiency in this treatment. Importantly, however, the two treatments should *not* differ in the second dimension of market efficiency—the level of unemployment. In particular, under the assumption of money-maximizing behavior of all players we should expect full employment also if contracts are not third party enforceable. To see this, note that—given our assumptions on production technology and payoff functions—it is more profitable for a firm to employ two workers at the minimum possible wage ($w = 0$) who exert minimum effort ($e = 1$) compared to employing just one worker at the same wage rate who exerts the same level of effort.

3.3.2 Fair-Minded Workers

There is mounting evidence both from laboratory (e.g., Fehr and Gächter 2000, Brown et al. 2004) and field studies (e.g., Bewley 1999, Cohn et al. 2008) which suggests that workers' performance in markets characterized by contractual incompleteness depends on whether they feel treated fairly by their employer (see also Akerlof 1982 and Akerlof and Yellen 1990 for an early theoretical assessment on how fairness preferences can impact labor market outcomes). In this Section, we will sketch how workers' fairness preferences could influence the results of our experiment. In particular, we will argue that the presence of fair-minded workers can lead to market outcomes where higher efforts are observed in the IC treatment, but

[3] Assuming that workers reject contract offers with net payoffs equal to their outside option yields $[w, \hat{e}] = [1, 1]$ in IC and $[w, \hat{e}] = [19, 10]$ in C.

where the level of unemployment is also higher compared to the subgame perfect equilibrium outlined in the previous paragraph.

Brown et al. (2004) have shown that the presence of fair-minded players in markets similar to our IC treatment gives firms incentives to build up long-term employment relationships and pay high wages to their workers. Workers in such long-term relations will adhere to the contract and provide non-minimal levels of work effort in most periods of the game. This outcome is profitable both for firms and workers: if there are sufficiently many fair-minded workers, it is profitable for firms to pay strictly positive rents even in the final period of the game, since fair workers reciprocate high wages by providing high efforts. If firms in addition re-employ only workers who adhere to the terms of the contract in pre-final periods (i.e., if firms use a policy of contingent contract renewal), *all* workers have an incentive to fulfill their contractually specified obligations in pre-final periods. Fair workers are willing to provide high levels of effort as long as the wage payment is considered fair. Selfish workers imitate fair workers because, by doing so, they have the chance to be re-employed and earn positive rents also in future periods.

In Appendix C.3 we show that similar, cooperative outcomes can be achieved in our IC treatment. However, they can involve above-minimal and involuntary unemployment. More specifically, if fair-minded workers have "strictly egalitarian" fairness preferences,[4] the fraction of fair-minded workers necessary to sustain cooperation in two-worker firms is higher than for firms which employ only one worker. Moreover, even if there are enough fair-minded workers, firms in our setup strictly prefer to employ only one worker when fairness requires that net surplus has to be shared equally between all members of a firm. The reason for the latter effect is the implemented production technology. Because of decreasing returns to scale, the second worker's marginal contribution to total surplus is relatively low. If workers' fairness concerns require the firm to share total surplus equally between the firm and all hired workers, it could be that employing two workers with fair contracts yields

[4]I.e., Workers employed in a one-worker firm expect the gains from trade to be shared equally between their firm and themselves, and workers employed in a two-worker firm expect that they receive one third of the (net) surplus.

lower total profits to the firm as compared to the situation where the firm employs only one worker with the corresponding rent-sharing contract. For the parameters specified in Section 3.2.3, this is the case (the derivation can be found in Appendix C.3). Employing only one worker and sharing rents equally with this worker is thus a dominant strategy for firms when the fraction of fair workers is high enough to sustain cooperation in one-worker firms. In this situation, firms in the IC treatment will ration jobs but pay high wages and succeed in eliciting above-minimal levels of effort. This, however, gives rise to involuntary unemployment.

Obviously, the relative profitability of employing one or two workers strongly depends on the wage payment necessary to induce work effort by fair-minded workers. Therefore, it is crucial for market outcomes what these workers consider to be a fair wage. So far we have assumed that fair-minded workers fulfill a contract if it offers strictly egalitarian rent-sharing between all members of a firm. This assumption is crucial for making job rationing profitable for firms. While the egalitarian fairness benchmark has received a lot of support in the literature (e.g., Fehr and Schmidt 1999, Fehr et al. 2007) and is also intuitively appealing for our setup, it is by no means obvious that it is the only plausible formulation of what constitutes fairness (see Konow 2003 for an overview of different fairness principles).

For instance, it is also possible that a fair-minded worker in a two-worker firm requires only that the surplus generated by *himself* is shared equally between the firm and him. This form of "relation-specific egalitarianism", where workers neglect the surplus produced by other workers, is consistent with evidence from multi-worker gift-exchange games. In experiments studying such games, it has been observed that some workers seem to care relatively little about the payoffs of co-workers and about the firm's payoff resulting from its interaction with these co-workers (see, e.g., Maximiano et al. 2007, Charness and Kuhn 2007). Under this weaker fairness benchmark, employing two workers is c.p. more profitable for firms since they can reap 50% of total surplus (instead of 33% under the strictly egalitarian fairness benchmark). Indeed, assuming the weaker fairness benchmark for the parameters specified in Section 3.2.3 changes the predictions for the IC treatment: if the fraction of fair-minded workers is high enough, employing two workers under "relation-specific" surplus sharing is *more* profitable for firms than rationing jobs and sharing surplus only with one worker (see Appendix C.3).

Which fairness benchmark describes workers' actual fairness preferences best is ultimately an empirical question. It is also possible that what constitutes a fair wage from a worker's perspective is itself influenced by market conditions. For instance, the prevailing unemployment level could influence workers' perception of what constitutes a fair wage. Akerlof (1982)'s efficiency wage model assumes such a negative relationship between the fair wage and the level of unemployment. Burks et al. (2007) have recently provided survey evidence supporting the view that labor market conditions might affect workers' fairness perceptions. If this is the case, we should expect that workers' willingness to provide effort for a given wage (and, therefore, also firms' profits) will depend on the rate of unemployment.

In the following section, we test the behavioral predictions with our experimental data. We first analyze whether the absence of third party contract enforcement in the IC treatment indeed gives rise to unemployment (Section 3.4.1). We then study market outcomes and behavior in the IC treatment in more detail. We investigate determinants of firms' hiring decisions and compare different contracting policies used by firms. Finally, we analyze how workers' behavior is influenced by market conditions.

3.4 Results

3.4.1 Unemployment

Our main interest concerns the level of unemployment, dependent on the degree to which firms can explicitly enforce workers' performance. Figure 3.1 depicts the level of "endogenous unemployment"[5] in the two treatments. Initially, unemployment

[5]Remember that, in both treatments, 3 workers were unemployed "by design" in every session and every period due to excess supply of labor. To measure "endogenous unemployment" we therefore calculate the total number of unemployed workers *minus* 3 and divide by the number of possible jobs (given that each of 7 firms could offer 2 vacancies in every period the number of possible jobs in each market is 14).

rates do not differ significantly between the different market environments (Mann-Whitney U-Test for period 1 observations only, $p = 0.217$)[6]. After the first few periods, however, we observe a strong increase in unemployment in the IC treatment, while the level of unemployment stays close to zero when contracts are third party enforceable. As a consequence, the overall level of unemployment differs substantially between the different market environments. When effort is not verifiable, the average unemployment rate is higher than 30% while it is only about 5% in the C treatment. The difference between treatments is highly significant (Mann-Whitney U-Test, $p < 0.01$). The strong difference in unemployment can also be seen when analyzing single experimental sessions: the highest average unemployment rate in a market where contracts are explicitly enforced is 8.7%. This number is still lower than the lowest average unemployment rate in one of the IC sessions (21.4%).

> **Result 1:** *We observe strong differences in unemployment between the two treatments. Under explicit contract enforcement (C treatment), unemployment levels are close to the minimal possible level. When effort is not verifiable (IC treatment), unemployment rises strongly before stabilizing on a relatively high level.*

The absence of third party contract enforcement has a strong and positive impact on the level of unemployment. We have hypothesized in Section 3.3 that unemployment in the IC treatment could be caused by firms who decide to ration jobs and pay fair, above market-clearing wages to those workers who are employed. In other words, the absence of explicit contract enforcement could lead to *involuntary* unemployment. Table 3.2 and Figure 3.2 provide evidence that this indeed the case. Table 3.2 summarizes the number of potential vacancies (Column 1), the number of firms' actual contract offers (Column 2), and the number of contract offers that are accepted by workers (Column 4) in the two treatments. In total, firms in each treatment could offer up to 1260 contracts. Given our assumptions on the production function, offering the maximum number of vacancies was efficient from a technological point of view in both treatments. However, the number of actual contract offers

[6] All non-parametric tests use session averages as independent observations. Reported p-values are always two-sided.

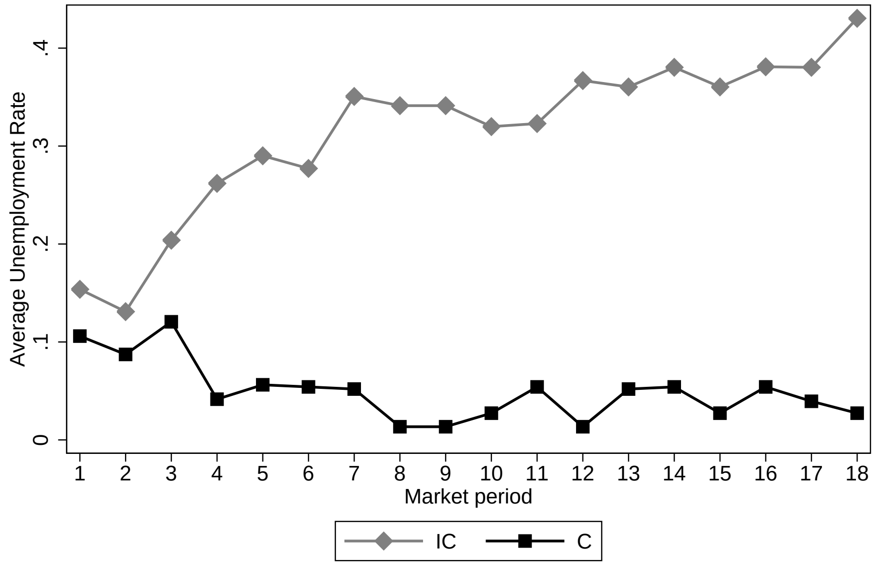

Figure 3.1: *Evolution of average unemployment over time*

differs strongly between treatments. While nearly 99% of potential vacancies (1242 out of 1260) are offered in the C treatment, firms in the IC treatment make only 68% (856 / 1260) of the possible contract offers. This indicates that the differences in unemployment between treatments are mainly caused by firms' decisions to ration jobs in IC rather than by differences in workers' contract acceptance between the two treatments. The number of accepted contract offers (Column 4) supports this impression: in both treatments, nearly all contract offers are accepted by some worker. This indicates that differences in workers' contract acceptance (which could, e.g., be caused by differences in workers' reservation wages) can not explain the differences in unemployment between treatments.

	Potential Vacancies	Actual Offers	Private Offers	Concluded Contracts
C treatment	1260	1242	207	1193
IC treatment	1260	856	616	849

Table 3.2: *Number of potential and actual (private) contract offers made by firms. Number of contracts accepted by workers.*

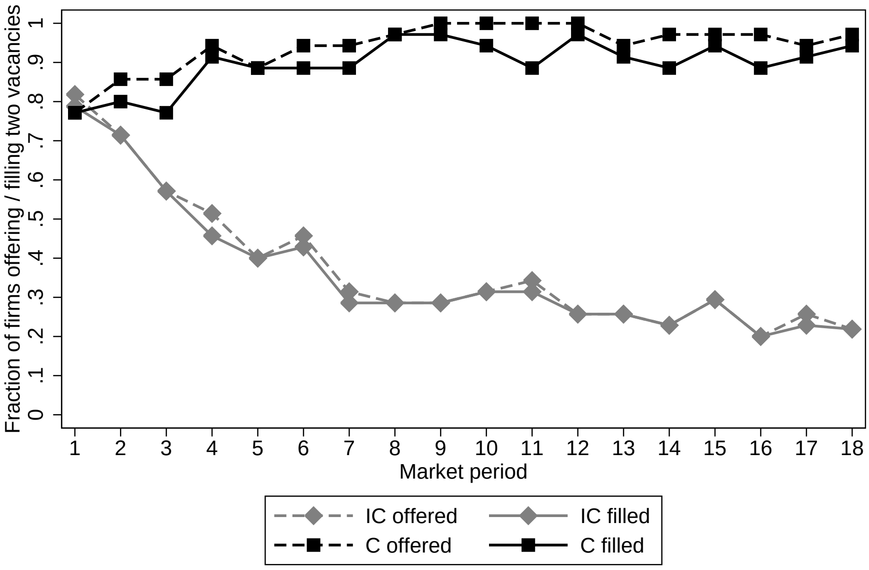

Figure 3.2: *Fraction of firms offering / filling two vacancies*

Figure 3.2 sheds more light on firms' hiring policy. The dashed lines depict the fractions of firms who offer a second vacancy in a given period. The solid lines compare the fractions of second contract offers which are accepted by workers. Paralleling the observation on the unemployment level, the fraction of firms who aim at hiring two workers is initially similar in both treatments. After a few periods, however, this fraction increases to more than 90% in the C treatment, but fewer and fewer firms decide to employ two workers in the IC treatment. In later periods of the experiment, a relatively stable fraction of merely 20–30% of firms offers two contracts in this treatment. The solid lines confirm that unfilled vacancies, which would be an indication for voluntary unemployment, are rarely observed in either treatment. If anything, the fraction of rejected contract offers is slightly higher in the C treatment.[7]

The prevalence of job rationing in the IC treatment also becomes apparent when we analyze the hiring decisions at the firm level. Table 3.3 reports the frequencies with which individual firms hire two workers. Paralleling the observations from

[7]The contract acceptance rate is 96% in the C treatment, compared to 99% in the IC treatment; see Table 3.2. Most of the rejected offers are contracts which offer zero or negative rents to the workers.

Figure 3.2, a vast majority of firms employ two workers in nearly all periods when work effort is verifiable. By contrast, in the IC treatment 74% of firms hire a second worker in less than half of the market periods. More than 50% of firms ration jobs even in more than 12 market periods.

Result 2: *Differences in unemployment are driven by firms' hiring decisions: when contracts are not explicitly enforceable, firms offer less contracts, and many firms employ only one worker in most periods of the experiment.*

Frequency of hiring two workers	0-2	3-5	6-8	9-11	12-14	15-18
Fraction of firms (C treatment)	–	–	–	–	0.171	0.829
Fraction of firms (IC treatment)	0.229	0.314	0.200	0.086	0.057	0.114

Table 3.3: *Percentage of firms who hire two workers in a given number of market periods (frequency with which individual firms employ two workers).*

3.4.2 Contracts in the Different Market Environments

The degree of third party enforceability apparently strongly affects overall market performance. In particular, the absence of explicit contract enforcement mechanisms leads to involuntary unemployment caused by firms' decision to ration jobs. In the remainder of this section, we analyze whether the channels through which unemployment is emerging in the IC treatment are also in line with our hypotheses. We first investigate whether the contracts which are most effective in eliciting work effort in this treatment differ from the profit-maximizing contracts in the C treatment. We then take a closer look at the IC treatment and analyze how the emergence of unemployment shapes the interaction between firms and workers.

We have argued in Section 3.3 that firms could try to elicit high work efforts by paying high wages and engaging in long-term employment relationships when third party contract enforcement is not feasible (IC treatment). On the other hand, the profit maximizing contract for firms in the C treatment should be the one which minimizes wage costs for a given—explicitly enforced—effort level. Furthermore,

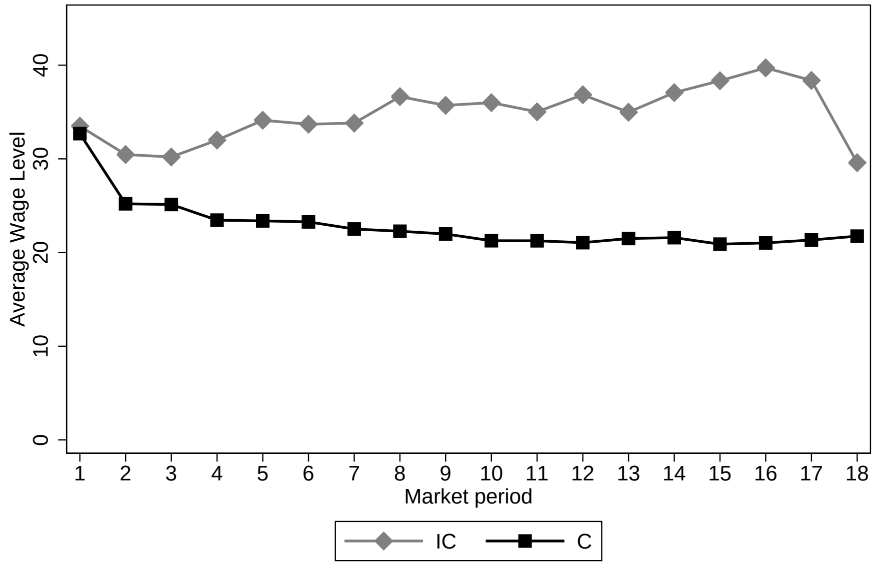

Figure 3.3: *Evolution of average wage over time*

there is no reason to (re-)employ specific workers via private contract offers when effort is verifiable.

Figure 3.3 and Table 3.4 show that these predictions are indeed borne out by our data. In Figure 3.3, we depict the average wage level in the two treatment conditions. When contracts are explicitly enforced, wages quickly converge to a level close to the market-clearing wage predicted in Section 3.3. Remember that in this treatment, profit-maximizing firms require the maximum possible effort of 10. Since unemployment benefits are set to zero, profit-maximizing firms should thus pay a wage equal to $c(10) = 18$. The wages observed in the C treatment are very close to this level: on average, firms pay 22.9 points to their workers (in the second half of the experiment, this value even decreases to 21.3 points). This contrasts sharply with the results for the IC treatment where effort is not verifiable. Firms in this treatment pay an average wage of 34.8 points to their workers. Moreover, wage differences across treatments increase over time. Overall, the differences in wages are statistically significant (Mann-Whitney U-Test, $p < 0.01$). Table 3.4 analyzes the profitability of these wage payments from a firm's perspective. In particular, we compare the profitability of paying higher wages and making private contract offers in the two treatments. As expected, paying higher wages when explicit contract

	Dependent variable: Firm profits					
	C treatment			IC treatment		
	(1)	(2)	(3)	(4)	(5)	(6)
Wage	-0.814***		-0.808***	0.485***		0.310**
	(0.137)		(0.143)	(0.085)		(0.101)
Private		-2.529	-0.818		18.250***	12.422***
		(3.619)	(2.787)		(2.828)	(2.565)
Constant	64.817***	46.671***	64.823***	8.288***	11.975***	5.420
	(3.609)	(0.802)	(3.541)	(2.808)	(1.293)	(3.006)
N	1193	1193	1193	849	849	849
R^2	0.241	0.007	0.242	0.150	0.160	0.204

Table 3.4: *Profitability of different contracts in the C treatment (Columns (1)–(3)) and the IC treatment (Columns (4)–(6)). The dependent variable is a firm's profit from a given contract. "Private" is a dummy equal to "1" ("0") if a firm-worker relation was initiated by a private (public) contract offer. *** indicates significance on the 1-percent level, ** significance on the 5-percent level, * significance on the 10-percent level. Reported standard errors (in parentheses) are adjusted for clustering on session level.*

enforcement is possible unambiguously decreases firms' profits (see Columns (1) and (3) of Table 3.4). Whether an employment relation is initiated by a private or public contract offer has no impact on profits in the C treatment (see Columns (2) and (3)). Therefore, it is not surprising that only few firms in this treatment make private contract offers (see Column 3 of Table 3.2).

These findings differ strongly to those obtained when explicit contract enforcement is not feasible where more than 70% of employment relations are initiated by a private contract offer (see Column 3 of Table 3.2). Firms indeed seem to care much more which worker they employ in the IC treatment. Is this selectivity in firms' hiring decision also reflected in a higher profitability of private contracts? Columns (5) and (6) of Table 3.4 confirm that this is the case. Trades initiated by a private contract offer are more profitable for firms. Moreover, treatments also differ strongly with respect to the relationship between firm profits and wages paid. As can be seen in Columns (4) and (6) of Table 3.4, paying higher wages increases firms' profits in the IC treatment. This indicates that the direct cost of paying higher wages is more than compensated by workers' increase in effort when generous wages are paid. Thus, the strong differences in wages observed in Figure 3.3 are well-founded from firms' perspective and reflect differences in the profitability of certain contracts in the two different market environments.

The differences in contracting between the two treatments is also reflected in the average duration of employment relationships. Since effort is verifiable, there is no need to build up long-term work relations in the C treatment; the average duration of a relationship in which a firm hires a specific worker (by directing a private contract offer towards him) is 1.35 periods. On the other hand, repeated interaction between firms and workers, i.e., the emergence of long-term work relationships is frequently observed in the IC treatment: conditional on hiring a worker through a private contract offer, firms employ the same worker on average for 2.90 periods. 51% of firms in this treatment rehire at least one specific worker for more than half of the market periods without interruption. The corresponding fraction of firms who does so in the C treatment is merely 3%.

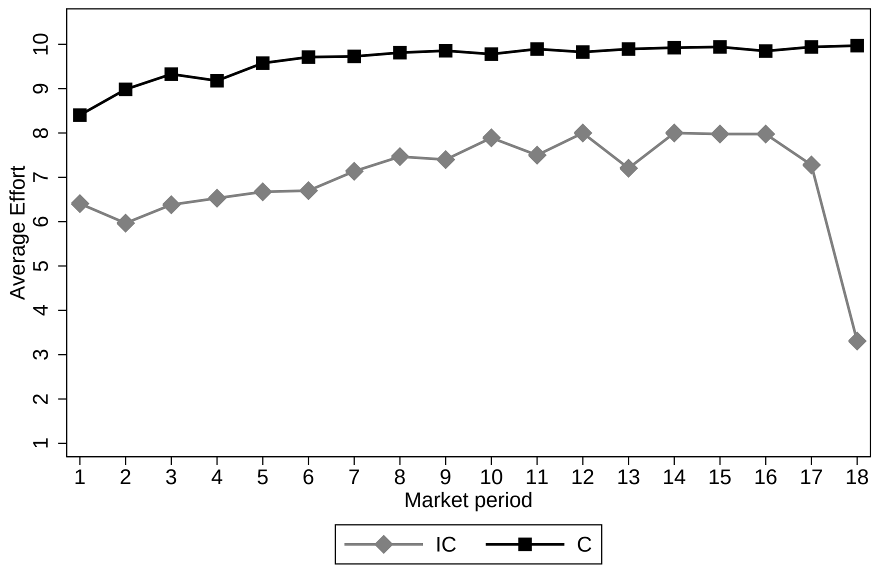

Figure 3.4: *Evolution of average effort over time*

Firms' strategy to ration job offers, pay high wages and repeatedly employ the same workers in the IC treatment gives powerful incentives for workers to provide effort, although effort is not explicitly enforced. Figure 3.4 depicts the average level of effort in the two treatments. In line with the theoretical hypotheses and previous evidence (Brown et al. 2004, Brown et al. 2008), the figure indicates that cooperation can indeed be sustained when third party contract enforcement is absent. While efforts in the IC treatment are lower than in the C treatment where they are explicitly enforced (Mann-Whitney U-Test, $p < 0.01$), the figure demonstrates that the implicit incentives provided through firms' contracting strategy are strong enough to elicit relatively high work effort. Although effort is not verifiable, workers choose an effort level of 8–10 in more than 50% of cases, and minimum effort is observed in only 12% of cases. Compared to the strong difference in unemployment between treatments, the difference in the second measure of market efficiency—workers' effort—is relatively small.

> **Result 3:** *Successful firms use different contracting policies in the two treatments. When effort is verifiable (C treatment), firms increase profits by employing workers at the lowest incentive compatible wage. When third party contract enforcement is absent, firms operate profitably if they pay generous wages and hire specific workers through private contract offers.*

3.4.3 Determinants and Consequences of Job Rationing

We now turn to the determinants of firms' decision to ration jobs in the IC treatment. We first concentrate on differences in profits between the different "types" of firms, i.e., firms who employ two workers and firms who ration jobs and employ only one worker.[8] Is the high unemployment in the IC treatment caused by overall higher profits of firms who employ only one worker, as suggested by the analysis of workers

[8] Strictly speaking, there are no stable "types" of firms in our experiment since firms could decide in each market period how many workers they want to hire. Although the correlation between individual firms' job rationing decisions in consecutive periods is highly positive ($\rho = 0.47$, $p < 0.01$), some firms also switch between hiring one or two workers. Firms keep their "type" in 75.6% of cases, two-worker firms become one-worker firms in 13.8% of cases, and one-worker firms employ two workers in the following period in 10.6% of cases. For ease of exposition, in what follows we will thus refer to a firm as "one-worker firm" ("two–worker firm") if it employs one (two) workers *in a given market period*.

with "strictly egalitarian" fairness preferences in Section 3.3? On average, one-worker firms in this treatment indeed earn somewhat higher profits. Average profits of firms who employ one worker in a given period are 36.9 points compared to 29.4 points in two-worker firms (see Column (1) in Table 3.5). However, this difference in profits is not statistically significant.

	Dependent variable: Firm profits					
	all market periods	periods 1–6	periods 7–12	periods 13–18	1-worker firms	2-worker firms
	(1)	(2)	(3)	(4)	(5)	(6)
2-worker	-7.499	-12.630***	-3.494	-0.987		
	(5.110)	(2.218)	(6.880)	(7.290)		
Period					0.370	1.087**
					(0.200)	(0.390)
Period18					-42.869***	-33.714**
					(7.141)	(11.625)
Constant	36.879***	36.587***	40.757***	33.395***	35.652***	22.283***
	(2.784)	(1.940)	(2.520)	(3.836)	(0.594)	(1.850)
N	623	208	209	206	397	226
R^2	0.030	0.102	0.008	<0.001	0.238	0.094

Table 3.5: *Profit differences between one-worker firms and two-worker firms in the IC treatment. The dependent variable is a firm's* total *period profit. "2-worker" is a dummy equal to "1" ("0") if a firm employs two (one) workers in a given period. "Period18" is a dummy equal to "1" for the final market period. *** indicates significance on the 1-percent level, ** significance on the 5-percent level, * significance on the 10-percent level. Reported standard errors (in parentheses) are adjusted for clustering on session level.*

The comparison of average profits over all market periods hides an interesting time pattern in profits for the different types of firms. Figure 3.5 depicts firms' profits in the IC treatment over time. The figure suggests that, initially, two-worker firms earn considerably less than one-worker firms, but that this difference in profits disappears over the course of the experiment. Comparing firms' profits in the first 6 periods yields a profit difference 12.6 points which is significant at the 1% level

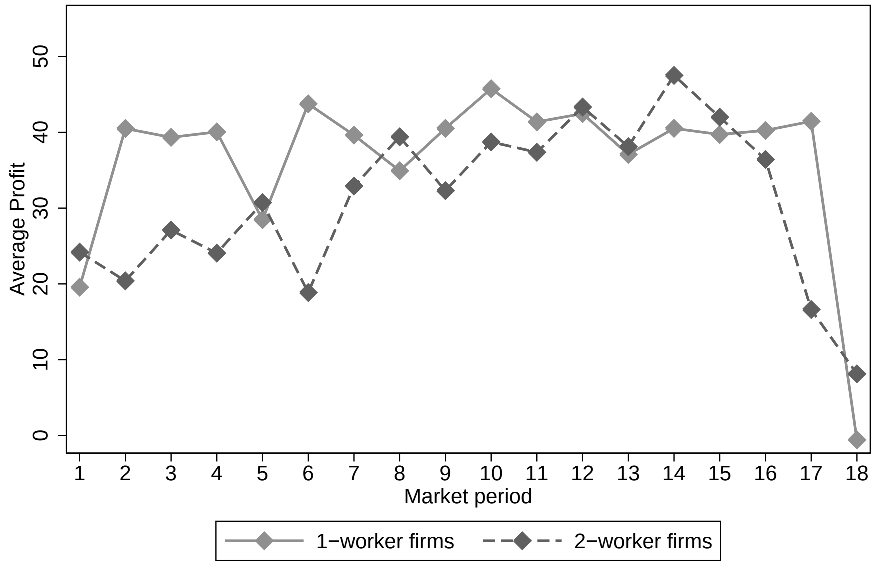

Figure 3.5: *Firm profits in 1-worker firms and 2-worker firms (IC treatment)*

(see Column (2) of Table 3.5). Paralleling the pattern observed in Figure 3.5, this difference in profitability between the different types of firms vanishes in later market periods (cf. Columns (3) and (4) of Table 3.5). Regressions analyzing the dynamics of profits for the two different types of firms separately indicate that the convergence in profits is caused by an increase in two-worker firms' profitability. Controlling for the drop in profits in the final period, two-worker firms significantly increase their profitability over the course of the experiment (see Column (6) of Table 3.5). At the same time, profits of one-worker firms increase weakly, but not significantly (see Column (5) of Table 3.5). This shows that a (stronger) increase of profits of two-worker firms rather than a decrease of profits in one-worker firms accounts for the convergence of profits between the different firm types.

> **Result 4:** *When third party contract enforcement is absent, one-worker firms are initially more profitable than two-worker firms. In later periods, profits in two-worker firms converge to the profit levels of one-worker firms.*

In Section 3.3, we have argued that the relative profitability of employing one or two workers in the IC treatment should depend on what constitutes a fair contract offer. If fair-minded workers have strictly egalitarian fairness preferences, one-worker firms will always be more profitable. The analysis of firm profits in the IC treatment has demonstrated that, in the beginning of the experiment, this seems to be case. In this sense, profit differences could account for the strong increase in unemployment in these periods. However, the analyses also have shown that a relatively stable fraction of firms employ two workers in later periods of the experiment (see Figure 3.2) and that these firms manage to earn similar profits as one-worker firms. The only way how two-worker firms can catch up in terms of profits are differences in the effort-wage relation between different types of firms. Because of the decreasing-returns-to-scale production technology, two-worker firms can operate as profitably as one-worker firms only if workers accept receiving a lower share of the surplus, and put in "extra" effort for a given wage level compared to workers in one-worker firms. Thus, if the effort-wage relation gets steeper over time in two-worker firms, this could explain the increase in these firms' profits observed in Figure 3.5.

Table 3.6 reports estimation results on the relationship between wages and workers' effort choices in the IC treatment. To avoid biases due to the strong drop in efforts in the final period of the market (cf. Figure 3.4), final-period observations are excluded from the analysis. The estimation shows that the effort-wage relations are positively sloped both for one-worker firms and two-worker firms (see Columns (1) and (4) of Table 3.6). However, the dynamics of the workers' effort choices indeed differ between the different types of firms. Controlling for wages paid, average effort in two-worker firms increases over time (see Column (5)). Column (6) confirms that this is the result of a steeper wage-effort relationship, i.e., for any given level of the wage workers in two-worker firms are willing to provide higher efforts over time. Columns (2) and (3) show that for workers employed in one-worker firms this effect is absent. For a given wage, these workers provide similar levels of effort throughout the experiment.

> **Result 5:** *The convergence in profitability of one-worker firms and two-worker firms can be attributed to an increasing steepness in the effort-wage relation of two-worker firms: for a given wage level, workers in two-worker firms are increasingly willing to provide high levels of effort.*

It seems as if the firms who employ two workers are able to use the increased

	Dependent variable: Effort					
	1-worker firms			2-worker firms		
	(1)	(2)	(3)	(4)	(5)	(6)
Wage	0.141***	0.141***	0.129***	0.164***	0.166***	0.134***
	(0.008)	(0.009)	(0.022)	(0.010)	(0.008)	(0.014)
Period		0.011	-0.037		0.082**	-0.051*
		(0.029)	(0.148)		(0.024)	(0.023)
Wage*Period			0.001			0.004**
			(0.003)			(0.001)
Constant	2.243**	2.157***	2.620**	1.518**	0.881***	1.857***
	(0.522)	(0.325)	(0.093)	(0.327)	(0.138)	(0.301)
N	372	372	372	438	438	438
R^2	0.643	0.644	0.644	0.595	0.610	0.620

Table 3.6: *Determinants of worker effort in the IC treatment. Estimations for 1-worker firms (Columns (1)–(3)) and 2-worker firms (Columns (4)–(6)). Observations from the final market period are excluded. *** indicates significance on the 1-percent level, ** significance on the 5-percent level, * significance on the 10-percent level. Reported standard errors (in parentheses) are adjusted for clustering on session level.*

level of unemployment in later periods to become more profitable. One potential explanation for this is that unemployment may serve as a worker disciplining device. Monetary payoffs for unemployed workers are zero in our experimental setup; therefore workers who are unemployed incur an immediate cost from not having a job. A consequence of the *increase* in unemployment is that being unemployed becomes relatively more costly. There are two reasons for this effect: first, due to firms' job rationing decisions, fewer and fewer job offers are available. Secondly, more unemployed workers compete for a job offer when unemployment is high. It should thus be less likely to get a job offer when unemployment is high.

Table 3.7 summarizes the job acquisition rate of previously unemployed workers in the IC treatment. The figures depict the fraction of workers who were unemployed in period $t-1$ and acquire a job in period t (either through a private or a public contract offer). The acquisition rate decreases from more than 50% in early periods to about 20-35% in later periods. The correlation between the (lagged) unemployment rate and the job acquisition rate is strongly negative ($\rho = -0.53$, $p < 0.01$).

Period (t)	2	3	4	5	6	7	8	9	10
Acquisition rate (t)	0.78	0.51	0.52	0.41	0.52	0.39	0.26	0.33	0.35
Unemployment ($t-1$)	0.16	0.14	0.21	0.27	0.30	0.29	0.37	0.36	0.36
Period (t)	11	12	13	14	15	16	17	18	
Acquisition rate (t)	0.35	0.28	0.32	0.25	0.22	0.17	0.20	0.32	
Unemployment ($t-1$)	0.34	0.34	0.37	0.37	0.39	0.37	0.40	0.39	

Table 3.7: *IC treatment: workers' probability of being hired in period t if unemployed in period* $t-1$.

Does the higher cost of being unemployed indeed discipline workers? Table 3.8 depicts results of probit estimations on the determinants of workers' shirking behavior (reporting marginal effects). The dependent variable is a dummy equal to "1" if a worker shirks, i.e., if he provides less than the contractually desired effort level. Explanatory variables include the wage stipulated in the contract, the type of contract offer, the current job acquisition rate, and a dummy equal to "1" if the worker has been unemployed in the previous period. In line with the previous results, the estimations indicate that workers who are hired via a private offer and who receive higher wages tend to shirk less. More importantly, however, the coefficients for the job acquisition rate and the unemployment dummy indicate a decrease in shirking rates through unemployment: controlling for the wage received, workers tend to shirk less when job offers are scarce (i.e., when the job acquisition rate is low) and when they have been unemployed before. Interestingly, we observe that workers who have experienced unemployment *themselves* shirk less once they find a job again. This suggests that, in our setup, unemployment disciplines especially the previously unemployed workers, and not necessarily all employed workers as in Shapiro and Stiglitz (1984).

Result 6: *The increase in unemployment in the IC treatment leads to a strong decrease in the job acquisition rate for unemployed workers. This makes being or becoming unemployed more costly. The higher pressure on workers is reflected in workers' contract fulfillment: workers shirk less when (i) jobs are scarce and (ii) when they have been unemployed previously.*

	Dependent variable: 1 if $e < \hat{e}$		
	(1)	(2)	(3)
Wage	-0.017***	-0.015***	-0.016***
	(0.003)	(0.003)	(0.003)
Unemployed before	-0.106***	-0.127***	-0.103**
	(0.036)	(0.025)	(0.040)
Private		-0.162**	
		(0.079)	
Acquisition rate			0.260*
			(0.148)
N	849	849	849
Pseudo R^2	0.162	0.172	0.174

Table 3.8: *Determinants of contract fulfillment in IC treatment. Probit estimations, reporting marginal effects. "Unemployed before" is a dummy equal to 1 if the respective worker had no contract in the previous period. "Acquisition rate" is the fraction of previously unemployed workers who find a new firm in a given period (see Table 3.7). *** indicates significance on the 1-percent level, ** significance on the 5-percent level, * significance on the 10-percent level. Reported standard errors (in parentheses) are adjusted for clustering on session level.*

We have seen that, for a given wage level, previously unemployed workers are more likely to fulfill a contract by providing the contractually agreed upon effort level. As a consequence, firms should be able to negotiate more advantageous contract terms, and appropriate a higher share of the surplus when unemployment is high. However, it is unclear so far why this helps especially the firms who employ two workers to become more profitable.

The answer lies in the differences in contract offers between one-worker firms and two-worker firms. Firms who employ two workers in a given period are more likely to hire workers who have been unemployed before. Although only roughly 20% of firms employ two workers in later periods, these firms account for almost two thirds (65.1%) of job offers to previously unemployed workers. One reason for this effect is that firms who employ two workers rely more on public job offers than one-worker firms. 70.7% of public contract offers come from firms who employ two workers.[9] Another reason lies in the policies of contract renewal employed by firms who employ one vs. two workers. Table 3.9 reports estimations on determinants of contract renewal in the IC treatment. Comparing the re-employment strategies of different firm types shows that, for a given effort level provided, workers in two-worker firms face a higher probability of being dismissed (see Column (1) of Table 3.9). The same holds even when we control for firms' total profits (see Column (2) of Table 3.9). Only if an *individual* worker's effort is so high that it yields firm profits comparable to those in one-worker firms, the likelihood of being re-employed does not differ from workers in one-worker firms (Column (3) of Table 3.9). Since workers in two-worker firms are less productive due to decreasing returns to scale, this implies a much tighter contract renewal policy of two-worker firms as compared to one-worker firms. In other words, for a given wage level a worker employed in a two-worker firm has to provide (costly) extra effort in order to attain re-employment chances comparable to a worker in a one-worker firm.

> **Result 7:** *Due to firms' strategies of offering and renewing work contracts, previously unemployed workers are more likely to be hired by firms who employ two workers.*

The higher likelihood of hiring workers from the pool of unemployed workers, together with the stricter contract renewal policy explains why two-worker firms more than proportionally profit from high unemployment in the IC treatment. Another consequence of the tighter re-employment policy of two-worker firms is a higher turnover rate and, overall, shorter employment relations compared to one-worker firms. Only in 19% of cases, two-worker firms hire the same workers in two consecutive periods. The comparable number for one-worker firms is 83%.

[9]This figure excludes contract offers from the first period where most firms (i) employ two workers and (ii) do so via public contract offers because they have not gained experience with specific workers.

	Dependent variable: 1 if worker is not re-hired		
	(1)	(2)	(3)
2-worker	0.165***	0.200**	-0.093
	(0.039)	(0.089)	(0.108)
Effort worker i	-0.138***		
	(0.012)		
Firm's total profit		-0.010***	
		(0.001)	
Firm's profit from worker i's effort			-0.016***
			(0.002)
Period	-0.022***	-0.019***	-0.023***
	(0.007)	(0.006)	(0.006)
N	807	807	807
Pseudo R^2	0.386	0.200	0.241

Table 3.9: *Determinants of worker dismissal. Probit estimations, reporting marginal effects. "2-worker" is a dummy equal to 1 if a firm employs 2 workers in a given period.* *** *indicates significance on the 1-percent level,* ** *significance on the 5-percent level,* * *significance on the 10-percent level. Reported standard errors (in parentheses) are adjusted for clustering on session level.*

Summarizing our findings, the outcomes in the IC treatment resemble an endogenously "segmented" labor market: on the one hand, the majority of firms employ only one worker when contracts are not third party enforceable. These firms are characterized by stable employment relationships based on private contract offers and contingent contract renewal. Moreover, these firms share the gains from trade generously, paying high wages to their workers which result in high work effort and high monetary payoffs both for workers and firms. On the other hand, a minority of firms operates successfully by employing two workers when unemployment is high. Interactions between these firms and their workers resembles those in "secondary" labor markets (Saint-Paul 1996). In comparison to one-worker firms, successful two-worker firms are characterized by low wages and thus low rent payments to workers. Moreover, these firms exhibit a higher turnover rate due to a tighter contract renewal policy and a higher fraction of public contract offers. High unemployment disproportionally helps these firms, since they are more likely to hire previously unemployed workers. For a given wage, these workers tend to shirk particularly little. The "secondary sector" firms are therefore able to pay lower wages to elicit a certain level of work effort. When unemployment is high, this strategy yields profits similar to those of firms who ration jobs and rely on long-term employment relations and high wages.

3.5 Concluding Remarks

In this chapter, we have analyzed the relationship between contract enforcement and the emergence of involuntary unemployment. In an experimental labor market where the trading parties can form long-term relationships, we compared a work environment where contracts were not third party enforceable to a situation where contracts were "complete" and effort was verifiable.

Our main findings are as follows: unemployment is much higher in the treatment without explicit contract enforcement. Importantly, unemployment in this treatment is involuntary, being caused by the firms' employment and contracting policy. Firms pay high wages but offer fewer vacancies than possible and technologically efficient. This policy, however, succeeds in eliciting high efforts from the employed workers. When complete contracts can be written, wages are close to the market clearing level, firms do not ration jobs, and unemployment is mostly voluntary.

To the best of our knowledge, our work is the first which empirically establishes a direct link between contract enforcement and the emergence of unemployment. This connection has long been discussed in efficiency-wage theories (Shapiro and Stiglitz 1984, MacLeod and Malcomson 1989, Akerlof and Yellen 1990). Our findings also contribute to the literature on the efficiency-wage foundation of dual labor markets (Bulow and Summers 1986, Saint-Paul 1996). Although firms face identical technological constraints in our setup, we observe a segmentation of the labor market when third party contract enforcement is not feasible. Most firms ration the number of jobs, and build up long-term employment relations which are characterized by generous rent-sharing between firms and workers. Behavior of these firms and their workers seems isolated from the market environment (e.g., the level of unemployment). On the other hand, there is also a minority of firms operating successfully without rationing job offers. These firms pay lower rents to workers, and use a tighter re-employment policy. Market conditions are very important for these firms: since they employ more workers through public contract offers and since previously unemployed workers shirk less, these firms become profitable only when unemployment is high.

Chapter 4

Behavior in Multi-Stage Elimination Tournaments

4.1 Introduction

Promotions play an important role for the provision of incentives in firms and other hierarchical organizations. According to Lazear and Gibbs (2008), they are probably the most important source of extrinsic motivation for middle managers in most firms. The prospect of being promoted to a better-paid job creates incentives to work hard, even if current income is not tied to performance. In virtually all firms those who get promoted compete again for subsequent promotions: in many companies, there are up to a dozen hierarchical levels between the CEO and entry-level management (Belzil and Bognanno 2008, Lazear and Gibbs 2008). The prevalence of internal labor markets is also underlined by the fact that hiring is concentrated on lower levels of the hierarchy and that positions on higher ranks are filled primarily through promotion (Baker et al. 1994, Bognanno 2001). In a seminal contribution, Rosen (1986) has modelled the competition for promotion in such hierarchies as a multi-stage elimination tournament where in each stage fewer agents are selected for the next step of the career ladder. Incentives generated in such tournaments depend on two important components of the organizational structure: the immediate wage increase for an agent who gets promoted and the option value of competing in further stages of the tournament, i.e., having the chance to earn even higher wages.

Although the importance of multi-stage elimination tournaments is undisputed, stringent empirical tests of their incentive effects are scarce. In this chapter, we provide a step towards closing this gap with the help of a controlled laboratory experiment. Our main questions are: do people take future stages of a tournament into account when deciding on current effort? Are multi-stage elimination setups behaviorally different from simple one-stage tournaments? How does the wage structure in multi-stage tournaments affect the provision of effort?

We study these questions by comparing three treatments. Our main treatment is a two-stage tournament (*TS*) in which four subjects compete for being promoted to the second stage. Promotion depends on subjects' output which is a function of costly effort and an individual noise term. The two subjects with the lowest output levels in the first stage are eliminated from further competition and receive a low wage. The two subjects with the highest output levels in the first stage are promoted, i.e., they are allowed to take part in the second stage where they compete against each other again. The subject with the highest second-stage output receives a high wage, whereas the other finalist is paid an intermediate wage. Parameters in this treatment are chosen to make the tournament incentive maintaining in the sense that equilibrium effort is identical in both stages (Rosen 1986).

We compare this treatment to a one-stage tournament (*OS*) in which four subjects compete once for two top positions. Wages for the promoted subjects are chosen such that the one-stage tournament is strategically equivalent to the first stage of *TS*. This means that promoted subjects in *OS* earn the sum of the intermediate wage in *TS* *and* the monetary equivalent of the second-stage option value in that tournament, implying that equilibrium effort is the same in both treatments. Comparing *OS* and *TS* thus allows testing whether strategic equivalence translates into behavioral equivalence. Our third treatment (*TSC*) is identical to the *TS* treatment with the exception that the wage structure is more convex, i.e., the intermediate wage is smaller in *TSC*. It is designed to study how subjects react to differences in the wage structure. In particular, we can test whether subjects—as predicted by theory—exert lower efforts in the first and higher efforts in the second stage under the more convex wage structure of *TSC* compared to *TS*.

Our findings can be summarized as follows: First, average behavior is remarkably close to the predictions of tournament theory in our one-stage treatment. This parallels findings of previous experiments on symmetric one-stage tournaments (e.g., Bull et al. 1987, Orrison et al. 2004). Second, behavior in the *TS* treatment indicates that subjects take the option value of future promotion possibilities into account when deciding on their work effort in multi-stage tournaments. Third, we also observe important departures from theoretical predictions in the *TS* treatment. Behavior in the first stage of *TS* differs strongly both from the one-stage treatment and from theoretical predictions. Subjects exert significantly higher efforts in the first stage of the two-stage tournament. Fourth, the results of the *TSC* treatment confirm the occurrence of excess effort provision in the first stage of the tournament. Subjects react only weakly to the change in the wage structure, implying that first-stage excess effort is even higher in *TSC*. Finally, the *TS* treatment is incentive maintaining in the sense that efforts in the second stage are not significantly different from first-stage efforts.

Most promotion tournaments that we observe in firms and other hierarchical organizations have multiple stages. Our results indicate that the mechanisms of incentive provision in multi-stage tournaments largely operate as suggested by theory. People do not only respond to differences in prizes, or wages, but are also motivated by the option value generated by future promotion possibilities. Our findings also provide insights with regard to the question whether one-stage tournaments are behaviorally equivalent to multi-stage designs. Adding one or more stages seems to make a fundamental difference, as people tend to exert excess effort in early stages of the tournament. This shows that one cannot necessarily draw inferences from simple one-stage setups to more complex tournaments.

The finding that people tend to exert excess effort in early stages of a multi-stage competition also has interesting organizational implications. If—as our findings suggest—multi-stage competitions provoke excess effort exertion by employees, this may offer a possible explanation why firms rely heavily on promotion based incentive schemes even if more direct means of performance assessment and compensation are available. Excess effort makes a promotion tournament "cheaper" for principals compared to other incentive schemes such as piece rates or bonus contracts. Moreover, excess effort influences the "optimal architecture" of promotion tournaments and other contests (see Moldovanu and Sela 2006).

The chapter complements previous studies which have used field data from executive compensation, sports, or agricultural production to evaluate predictions of tournament theory (e.g., Ehrenberg and Bognanno 1990, Knoeber and Thurman 1994, Bognanno 2001). It has been found that many implications of tournament theory are consistent with empirical observations: for instance, higher prizes tend to increase performance (Ehrenberg and Bognanno 1990), wage profiles in the field are convexly increasing (Eriksson 1999), and winner prizes increase with the number of competitors (Bognanno 2001). An advantage of testing tournament theory with laboratory experiments is the possibility to directly test *pure incentive effects* of tournaments. For instance, one can measure agents' behavioral reactions to changes in the wage structure without being concerned that the composition of the workforce might change due to the modified incentive scheme (Lazear 2000, Dohmen and Falk 2006). Experimental data also allows to rule out influences of unobservable variables that might affect actual promotion decisions in the field, like agents' soft skills or supervisor favoritism.

A number of previous papers have studied various aspects of tournaments in experiments, such as the effects of different prize spreads (Bull et al. 1987, Harbring and Irlenbusch 2003), sabotage activities (Falk, Fehr and Huffman 2008, Harbring and Irlenbusch 2008) or asymmetries in promotion chances (Schotter and Weigelt 1992). In contrast to our work, this literature concentrates on one-stage tournaments. In view of our findings it is not clear to what extent these previous findings translate to multi-stage setups. For instance, more competitive behavior (excess effort) in early stages might also lead to an increase in sabotage activities in multi-stage tournaments.

The remainder of the chapter is organized as follows: the next section presents a simple model of multi-stage elimination tournaments on which our experiment is based. Section 4.3 discusses our experimental design and derives hypotheses. Section 4.4 shows the main results and Section 4.5 concludes.

4.2 A Simple Model of Multi-Stage Elimination Tournaments

We consider a simple elimination tournament in which four identical agents compete for promotion.[1] The promotion decision depends on relative output produced by the agents. Competition consists of two stages: in the first stage, all four agents compete against each other. The two agents with the lowest output levels in that stage receive a wage w_{low} and are eliminated from further competition. The two agents with the highest output levels in the first stage are promoted, i.e., they are allowed to take part in the second stage (or "final") where they compete against each other again. The agent who produces more output in the second stage receives a wage w_{high}, whereas the other finalist gets an intermediate wage w_{med}. Note that the decision of who receives w_{high} or w_{med} does not depend on the first-stage output of the finalists.

This two-stage elimination tournament can be modeled as follows. In the first stage of the tournament four agents $i = 1, 2, 3, 4$ compete against each other. Agents who participate in stage $k \in \{1, 2\}$ individually produce output $y_{i,k}$ according to the production function

$$y_{i,k} = e_{i,k} + \epsilon_{i,k}$$

where $e_{i,k}$ denotes the effort level that agent i exerts in stage k. $\epsilon_{i,k}$ is a random shock faced by agent i in stage k. Shocks are assumed to be drawn independently for each agent in each stage. For simplicity, we assume that $\epsilon_{i,k}$ is uniformly distributed on the interval $[-q, q]$.[2] Agent i's output in stage k does not depend on previous effort or output and the production technology is identical for all agents i in all stages k. Agents bear the cost of effort exertion. We assume the cost function to be

[1] Most of the assumptions below follow the classic (one-stage) tournament model introduced by Lazear and Rosen (1981).

[2] Virtually all tournament experiments use the uniform distribution, primarily because its concept is easy to understand for experimental subjects. The predictions of the model, however, can be generalized to other distributions of shocks. Lazear and Rosen (1981) discuss which conditions have to be fulfilled for the existence of an equilibrium in pure strategies.

of the following form:

$$C(e_{i,1}, e_{i,2}) = \frac{e_{i,1}^2}{c} + \frac{e_{i,2}^2}{c}$$

Note that this specification implies separability of costs across stages, i.e., in line with Rosen (1986) there is also no carry-over of costs between stages. Furthermore, we assume that agents are identical and risk-neutral with utility functions which are additively separable in wages and effort costs

$$U_i(w, e_{i,1}, e_{i,2}) = w - C(e_{i,1}, e_{i,2})$$

For the derivation of equilibrium predictions we restrict our attention to the set of symmetric subgame perfect Nash equilibria. The two-stage tournament can be solved by backward induction. Because (i) the decision who wins the second stage solely depends on the output of the finalists in this stage, (ii) there is no cost carry-over between stages and (iii) the random terms are independently distributed both across stages and agents, the final of our two-stage tournament is equivalent to a simple one-stage tournament in which two participants compete for a promotion. Given that agent i has reached the second stage where two agents compete for one winner prize w_{high} and one loser prize w_{med}, he chooses stage-two effort $e_{i,2}$ in order to maximize an expected utility function of the following form:[3]

$$EU_i(w_{high}, w_{med}, e_{i,2}, e_{j,2}) = \pi(y_{i,2} > y_{j,2})w_{high} + [1 - \pi(y_{i,2} > y_{j,2})]w_{med} - C(e_{i,1}, e_{i,2})$$

$\pi(y_{i,2} > y_{j,2})$ denotes the probability that i's output in stage two is greater than the output of agent j. With our assumptions regarding the production function and random terms, this expression can be rewritten as follows:

$$EU_i(w_{high}, w_{med}, e_{i,2}, e_{j,2}) = F_{\epsilon_{j,2}-\epsilon_{i,2}}[e_{i,2} - e_{j,2}](w_{high} - w_{med}) + w_{med} - C(e_{i,1}, e_{i,2})$$

where $F_{\epsilon_{j,2}-\epsilon_{i,2}}[\cdot]$ denotes the cdf of the difference between random terms $\epsilon_{j,2}$, $\epsilon_{i,2}$. Maximizing $EU_i(\cdot)$ over $e_{i,2}$ yields the following first-order condition:

$$f_{\epsilon_{j,2}-\epsilon_{i,2}}(e_{i,2} - e_{j,2})(w_{high} - w_{med}) = \frac{\partial C(\cdot)}{\partial e_{i,2}}$$

[3] The parameters chosen in our experimental treatments (see below) ensure that all participation constraints are fulfilled. We therefore do not explicitly consider them here.

Assuming symmetry yields $f_{\epsilon_{j,k}-\epsilon_{i,k}}(0) = \frac{1}{2q}$ for $\epsilon_{i,k}, \epsilon_{j,k} \sim U[-q,q]$. The symmetric subgame perfect Nash equilibrium of the two-stage tournament thus entails the following second-stage effort level $e_{i,2}^{*,TS}$ for the two agents who participate in the final:

$$e_{i,2}^{*,TS} = \frac{(w_{high} - w_{med})c}{4q}$$

Given that both finalists play this equilibrium, the expected utility gain in the final, i.e., the continuation value for an agent in the first stage is given as follows

$$EV_{i,2} = w_{med} + \frac{1}{2}[w_{high} - w_{med}] - \frac{\left(e_{i,2}^{*,TS}\right)^2}{c}$$

An agent who reaches the final earns a wage of w_{med} for sure. By exerting stage-two effort $e_{i,2}^{*,TS}$, he has the chance to receive the higher wage w_{high} instead. In the symmetric equilibrium, this occurs with probability $\frac{1}{2}$. Moreover, he has to pay the cost of effort exertion in the second stage.

Turning to the analysis of the tournament's first stage, it is obvious that (in expected values) this stage can be modeled as a one-stage tournament between four agents with two winner prizes $EV_{i,2}$ and two loser prizes w_{low} (see Rosen 1986). The derivation of equilibrium effort for such a tournament follows the same steps as above. Alternatively we can apply a result from Orrison et al. (2004) who show that equilibria of fully symmetric one-stage tournaments are not affected by "organizational replication" for our specification of the production function, cost function, and random terms. This implies that an equilibrium in a tournament with two identical participants and one winner prize is also an equilibrium in a tournament with four identical participants and two winner prizes.[4]

As a shortcut we can therefore use the solution for $e_{i,2}^{*,TS}$ and simply replace w_{high} and w_{med} with $EV_{i,2}$ and w_{low} to obtain the equilibrium effort level for the first stage:

$$e_{i,1}^{*,TS} = \frac{(EV_{i,2} - w_{low})c}{4q} = \frac{\left(w_{med} - w_{low} + \frac{1}{2}\left[w_{high} - w_{med}\right] - \frac{\left(e_{i,2}^{*,TS}\right)^2}{c}\right)c}{4q}$$

[4] More generally Orrison et al. (2004) show that an equilibrium in a tournament with n participants and $\frac{1}{2}n$ winner prizes is still an equilibrium in a tournament with mn participants and $\frac{m}{2}n$ winner prizes.

This expression illustrates the two components of incentive provision in multi-stage tournaments. By winning the first stage and qualifying for the final, an agent receives an *immediate wage gain* $(w_{med} - w_{low})$, but additionally has the option to compete in the final and win the top prize w_{high}. The *value of this option* is

$$\frac{1}{2}[w_{high} - w_{med}] - \frac{\left(e_{i,2}^{*,TS}\right)^2}{c}.$$

Several aspects of the model deserve special emphasis. First, our design closely follows the original model of elimination tournaments by Rosen (1986) with one notable exception: instead of having two semifinals with two participants each, who compete for one slot in the final, we analyze a setup with four participants competing for two slots in the final. While both variants are theoretically equivalent for symmetric agents, we employ the latter because it allows us to design a one-stage tournament which is procedurally as close as possible to our main treatment.[5] Note also that we abstract from heterogeneous abilities of participants in order to keep the design as simple and parsimonious as possible. This allows us to focus on the incentive aspect of elimination tournaments by ruling out selection of more able individuals into higher positions of a hierarchy. Including the selection aspect of tournaments would be an interesting follow-up to our study.

4.3 Experimental Design

4.3.1 Treatments and Hypotheses

Our experiment comprises of three treatments that allow us to study behavior in multi-stage tournaments from different angles. For all treatments our benchmark is the prediction of the symmetric subgame perfect Nash equilibrium. The main treatment *TS* is a two-stage elimination tournament with four participants competing for w_{high}, w_{med}, and w_{low} as discussed in the previous section. We compare this treatment to a one-stage tournament (*OS*) in which four subjects compete for two

[5] In particular—as will become clear in the next section—both tournaments have the same number of participants and subjects compete for the same number of promotions.

top positions. The two subjects with the highest output levels receive a wage w'_{med} in *OS* while the two losers of the competition receive a wage w_{low}. The *OS* treatment fulfills several purposes. First, it serves as a validity check for our results given that a number of studies on one-stage tournaments already exists. In particular, our parametrization of this treatment is very close to a treatment from Orrison et al. (2004).

More importantly, however, the *OS* treatment allows us to investigate whether one-stage tournaments are behaviorally different from multi-stage ones. To investigate this question we design *OS* such that it is strategically equivalent to the first stage of the two-stage tournament *TS*. As discussed in the previous section the first stage of a two-stage tournament can be interpreted as a one-stage tournament in which agents compete for the expected value of participating in the second stage. Strategic equivalence between *OS* and the first stage of *TS* is thus achieved by keeping w_{low} constant and choosing:

$$w'_{med} = EV_{i,2} = w_{med} + \frac{1}{2}[w_{high} - w_{med}] - \frac{\left(e_{i,2}^{*,TS}\right)^2}{c}$$

In other words, the wage for the promoted agents in the one-stage tournament (w'_{med}) is equivalent to the wage w_{med} from *TS plus* the option value of participating in the final of *TS*. This choice implies that equilibrium effort levels in the *OS* treatment and in the first stage of the *TS* treatment are the same. We can therefore formulate the *behavioral equivalence hypothesis*:

$$e_{i,1}^{OS} \quad = \quad e_{i,1}^{TS}$$

In addition to comparing behavior in tournaments with different numbers of stages, we are interested in how a change in the wage structure influences behavior in multi-stage tournaments. This comparative statics exercise is of practical interest. Wages are one of the variables which (within certain bounds) are most amenable to manipulation in organizational design. We therefore compare our main treatment to a second two-stage tournament with a more convex wage profile (*TSC*). It is identical to the *TS* treatment with the exception that the intermediate wage w''_{med} in this treatment is smaller than in *TS*. This implies weaker incentives (lower equilibrium effort) in the first stage and a higher wage spread and higher equilibrium effort in the second stage (*wage structure hypothesis*):

$$e_{i,1}^{TSC} < e_{i,1}^{TS}$$
$$e_{i,2}^{TSC} > e_{i,2}^{TS}$$

Finally, we analyze behavior in our main treatment across stages. In the *TS* treatment wages are chosen such that equilibrium efforts in the first and second stage are equal. Remember that the two elements of incentives in multi-stage tournaments are the wage spread and the option value of competing for further promotions. In the final stage, the option value is zero because there are no further promotions beyond that stage. To make the tournament in the *TS* treatment incentive maintaining in the sense of Rosen (1986), this decrease in the option value in the second stage is offset by an appropriate increase in the wage spread ($w_{high} - w_{med} > w_{med} - w_{low}$). Comparing behavior across stages thus allows to test the *incentive maintenance hypothesis*:

$$e_{i,1}^{TS} = e_{i,2}^{TS}$$

Experimental parameters and the resulting equilibrium efforts for all treatments are shown in Table 4.1. When deciding on their efforts, subjects could choose any integer $e_{i,k} \in \{0, 1, ..., 125\}$. The effort costs in each stage are given by

$$C(e_{i,k}) = \frac{e_{i,k}^2}{2250}$$

and $\epsilon_{i,k} \sim U[-60, 60]$.[6] The parameters chosen imply equilibrium efforts of 74 in both stages of *TS* and in *OS*, the lower intermediate wage in *TSC* changes equilibrium efforts to 42 in the first stage and 100 in the second stage of this treatment.

4.3.2 Experimental Procedures

The experiment was conducted at the BonnEconLab of the University of Bonn. A total of 96 subjects in six sessions were divided into groups of four. We employed

[6]Note that, while in equilibrium all players make positive profits, the range of feasible efforts and the specification of the cost function imply that in principle subjects could make losses. In the few cases where this occurred, losses were deducted from the showup-fee.

Treatment	*OS*	*TS*	*TSC*
c	2250	2250	2250
q	60	60	60
w_{high}	–	20	20
w'_{med} / w_{med} / w''_{med}	13.62	12.11	9.33
w_{low}	5.73	5.73	5.73
$e^*_{i,1}$	74	74	42
$e^*_{i,2}$	–	74	100

Table 4.1: *Experimental parameters and resulting equilibrium predictions.*

a one-shot between-subjects design, i.e., subjects participated either in the *OS*, the *TS* or the *TSC* treatment. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

Before the tournament started, subjects received detailed written instructions on the respective treatment they took part in (*OS*, *TS* or *TSC*). These were neutrally framed and did not contain potentially value-laden terms like "tournament", "final", "winner", etc.[7] After reading the instructions subjects completed several control questions. The experiment started only after all participants had answered all control questions correctly. During the tournament, subjects simultaneously entered their effort decision and were then asked to state their expectations about the other participants' efforts on the next screen. This question was not announced beforehand. After the first stage participants in *TS* and *TSC* were only informed about the realization of their own random draw and about whether they had been promoted to the second stage. The finalists then again made an effort choice and entered their expectation about their opponent's effort. At the end of the tournament subjects in all treatments were informed about their earnings and asked to fill in a questionnaire. The structure of the experimental session ensured that subjects' decisions in the tournament can be treated as independent observations. The whole experimental session lasted on average 100 minutes and subjects earned an average of 18.25 Euro (1 Euro = 1.26 USD at the time of experiment), including a showup-fee

[7]Instructions for the *TS* treatment and the schedule of effort costs handed out to participants can be found in Appendix D.5 and Appendix D.6, respectively.

of 4 Euro and a fixed payment of 3 Euro for completing the questionnaire.

Note that our experimental procedures differ from previous tournament experiments in that we implement a one-shot interaction structure. Previous experiments have typically used repeated interactions. The advantage of the latter is that it allows for learning, which is potentially important given the non-trivial decision environment in tournaments. A potential downside, however, is that repeated game structures question the validity of static equilibrium predictions. Since we are explicitly interested in testing theoretical predictions, we decided to use a one-shot design. This has the additional advantage that stakes in the one-shot interaction are relatively high. We check the regularity of our results by comparing the outcome of our *OS* treatment with those of similar repeated tournaments. Finding similar results would make us confident that our main findings are not driven by the one-shot character of our set-up, but instead by treatment differences.

4.4 Results

In this section we first test whether the results for the one-stage tournament replicate earlier findings from similar tournaments. We then study the dynamic aspect of multi-stage tournaments by comparing the one-stage tournament *OS* to the strategically equivalent first stage of the two-stage tournament *TS*. In a third step, we investigate the effects of differences in the wage structure on effort provision by comparing the two-stage tournaments *TS* and *TSC*. Finally, we address the question whether the *TS* treatment is incentive maintaining by analyzing behavior in the first and second stage of the tournament.

4.4.1 Behavior in the One-Stage Tournament

Table 4.2 reports effort decisions in the *OS* treatment (Column 1). Two points are worth noting. First, efforts are on average very close to the theoretical predictions. While the average effort of 69.9 is slightly below the Nash prediction of 74, median effort coincides exactly with the predicted effort level. Second, there is substantial heterogeneity in subjects' behavior (see also Figure 4.1).

Both observations are in line with previous findings from symmetric one-stage designs (e.g., Bull et al. 1987, Eriksson et al. 2006). In particular, our results replicate those found by Orrison et al. (2004) who observe an average effort of 73.3 for an almost identical tournament which was repeated 20 times using lower stakes. The similarity of our results to those of Orrison et al. (2004) show that one of the most important findings in the experimental literature on symmetric promotion tournaments—average effort being close to Nash predictions—is quite robust with respect to using one-shot vs. repeated interactions and with respect to increased stake sizes.

Result 1: *Average behavior in the one-stage tournament is close to the predictions of the symmetric Nash equilibrium.*

4.4.2 Testing Behavioral Equivalence

Our one-stage tournament and the first stage of the two-stage tournament *TS* are strategically equivalent in the sense that the wage w'_{med} in the one-stage tournament includes the equilibrium option value of participating in the second stage of the two-stage tournament. A comparison of $e_{i,1}^{OS}$ and $e_{i,1}^{TS}$ therefore serves as a test of how subjects in the two-stage tournament perceive this option value. If subjects, for instance, do not take the option value into account when deciding on their first stage effort in *TS*, efforts in this stage should be lower compared to the *OS* treatment. If subjects evaluate the option value correctly, efforts in both treatments should be identical. Columns 1 and 2 of Table 4.2 show that behavior differs strongly between the two treatments. Average effort in the first stage of the *TS* treatment is 89.2, while median effort is 91. Thus, subjects behave much more competitively in the multi-stage tournament, exerting efforts which are more than 20% higher than those of their counterparts in the *OS* treatment. A Mann-Whitney U-test confirms that this treatment difference is highly significant ($p = 0.005$, two-sided).

Comparing effort levels in the two treatments to the theoretical predictions derived in Section 4.3 indicates that it is excess effort in *TS* rather than "too low"

Treatment	*OS*	*TS*	*TSC*
Average Effort	69.9	89.2	82.4
Median Effort	74	91	83
min(Effort)	1	40	40
max(Effort)	125	125	125
Variance	913.3	359.7	605.4
e^*	74	74	42

Table 4.2: *First stage behavior in the* OS, TS, *and* TSC *treatment.*

effort in *OS* that drives the treatment difference. A t-test with the null hypothesis that efforts are equal to Nash predictions confirms this: the null hypothesis is rejected in the *TS* treatment ($p < 0.001$), but cannot be rejected in *OS* ($p = 0.453$).[8] Subjects' effort choices suggest that they are not naive in the sense that they ignore the second stage. Quite to the contrary, the two-stage elimination tournament seems to trigger especially competitive behavior in the first stage.

The treatment difference is not just driven by some subjects choosing extreme effort levels in the *TS* treatment. A closer look at the distributions of first stage efforts in Figure 4.1 reveals instead that the whole effort distribution is shifted to the right in the *TS* treatment. As a consequence, efforts are less dispersed (Levene's test for equality of variances, $p < 0.05$, two-sided). The effort distribution illustrates that exerting excess effort is quite widespread in *TS*: 84% of subjects choose efforts higher than the equilibrium effort level of 74. This compares to only 47% in the *OS* treatment. The strong difference between treatments is also reflected at the lower tail of the distribution. While the lowest effort in the *OS* treatment is 1, no subject

[8]Note that the derivation of equilibrium predictions is based on the assumption of risk neutrality. We check the validity of this assumption by eliciting subjects' risk preferences with an incentive compatible lottery procedure (a detailed description of the procedure can be found in Appendix D.7). It turns out that our experimental subject pool is close to risk neutrality: the median subject in all three treatments is risk neutral, and the certainty equivalent of more than 50% of subjects lies in a range of $+/-0.25$ Euro around the risk neutral certainty equivalent (equal to 2 Euro for the lottery that was used). In addition, effort levels in *TS* are above theoretical predictions irrespective of the subjects' degree of risk aversion. Thus excess effort is not driven by subjects' risk preferences.

exerts effort below 40 in the *TS* treatment.

> **Result 2:** *Efforts in the first round of* TS *are significantly higher than in the* OS *treatment. This difference is driven by excess effort in* TS.

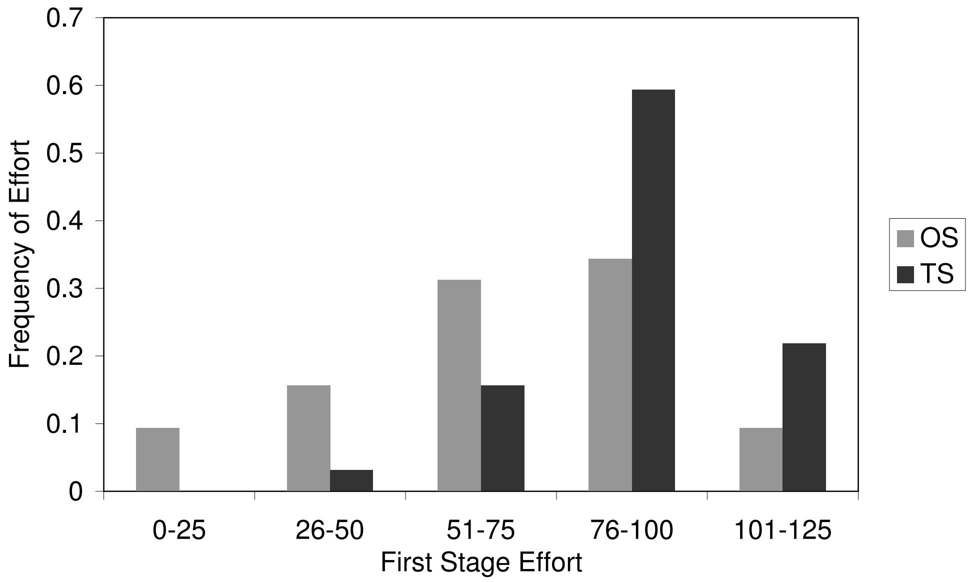

Figure 4.1: *Frequency of effort choices in the* OS *treatment and the first stage of the* TS *treatment.*

4.4.3 Wage Structures in Two-Stage Tournaments

In our next treatment comparison we investigate the behavioral effects of different wage structures in two-stage tournaments. Our main interest concerns the question whether first stage excess effort is a robust phenomenon. Is excess effort in the first stage an artefact of the specific wage structure we used in the *TS* treatment, e.g., is it specific to the incentive maintaining wage spreads? Or do we observe excess effort even if we reduce the incentive to provide effort in the first stage? To analyze these questions we implemented the *TSC* treatment, which is identical to the *TS* treatment except that the subject who produces less output in the second stage receives only $w''_{med} = 9.33$ Euro instead of $w_{med} = 12.11$ Euro. This more convex wage profile has the following theoretical implications. Incentives to provide effort in the first stage are weakened—the equilibrium effort level in the first stage of *TSC* is only 42 instead of 74 in the *TS* treatment. In the second stage, equilibrium effort increases from 74 to 100 (see Table 4.1).

Column 3 of Table 4.2 summarizes behavior in the first stage of the *TSC* treatment. Efforts in the first stage are much higher than theoretically predicted. The average effort level is 82.4 points, about 40 points higher than the equilibrium effort level of 42. A t-test with the null hypothesis that efforts are equal to the Nash prediction confirms that effort choices are significantly above the equilibrium prediction ($p < 0.001$). Indeed, about 88% of subjects choose efforts higher than 42. This suggests that excess effort in the first stage of our two-stage tournaments is robust to differences in the wage structure. Comparing effort choices in the *TSC* treatment to those observed in *TS* reveals that subjects react only weakly to the change in the wage structure between the two treatments. Average effort in *TSC* is only 7 points lower than in *TS*. This finding is striking given that there is a 32-point difference in equilibrium efforts. Despite the fact that first stage incentives are much weaker than in *TS*, effort choices do not differ significantly between the two treatments (Mann-Whitney U-test, $p = 0.245$, two-sided).

> **Result 3:** *The more convex wage structure in* TSC *induces even higher excess effort in the first stage. First stage efforts in* TS *and* TSC *are not significantly different, despite much weaker incentives in the* TSC *treatment.*

4.4.4 Testing Incentive Maintenance

Finally, we turn to behavior in the second stage of the *TS* treatment. Remember that parameters were chosen such that the tournament in *TS* is incentive maintaining, i.e., equilibrium effort levels are the same (equal to 74) in both stages. We know already that efforts are above the equilibrium prediction in the first stage of *TS*. In this sense, we can reject the hypothesis of *equilibrium* effort choices in both stage of *TS*. It remains to show, however, whether effort levels are the same in both stages of the tournament, or whether they differ. Does the two-stage character of *TS* induce above equilibrium effort also in the second stage or do players reduce efforts relative to their first stage behavior?

In turns out that the the evidence is mixed. On average, effort decreases when comparing behavior across stages in the *TS* treatment. Average effort in the first stage is 89, in the final it goes down to 82. The median effort choice decreases from 91 to 87.5. Note that average effort in the first stage includes the efforts of those who did not make it to the second stage. Since—by design of the promotion tournament—the latter usually exerted lower effort, the decrease from stage one to stage two is larger if we consider only finalists' behavior. Their average effort in the first stage is 96, implying that on average finalists decrease their effort by 14 points. However, these numbers hide considerable heterogeneity on an individual level: the fraction of finalists who decrease their effort in the second stage is only slightly higher than the fraction of subjects who increase it (50% and 44%, respectively). It is therefore not surprising that the overall decrease in finalists' efforts is insignificant (Wilcoxon signed rank test, p-value = 0.289). However, those who adjust their effort downwards on average do so much stronger than those who raise their effort.

Does this mean that excess effort in the first stage just mirrors the lower efforts in the second stage? That is, do subjects *expect* low effort levels and thus low effort costs in the second stage and therefore increase their first stage effort due to a higher (perceived) option value? While this might be the case for some finalists, we can rule out that it accounts for first stage behavior on a more general level as, on average, effort choices in the second stage are still above the equilibrium prediction. We can also use the observed first stage effort choice of a subject to calculate the option value implicitly underlying her decision. Using this option value, we can then construct the (hypothetical) second stage effort level which would rationalize the first stage effort choice of the subject at hand. For instance, if a subject beliefs that all players will exert zero effort in the second stage, her subjective option value increases to $\frac{1}{2}[w_{high} - w_{med}]$ since $\frac{\left(e_{i,2}^{TS}\right)^2}{c} = 0$ (see Section 4.2). Assuming this option value instead of the equilibrium option value rationalizes a first stage effort of 97 (using the wages and cost parameters of *TS*). Conducting the calculation for the first stage effort choice of the median subject in *TS* (equal to 91) yields an implied second stage effort of 37. I.e., if the median subject had expected a second stage effort level of 37, the perceived option value would rationalize her observed first stage effort choice. This value is, however, far below the actual effort choices in the second stage (cp. Table 4.2).[9] The expectation of low second stage effort levels can thus not account for the observed excess effort in the first stage of *TS*.

> **Result 4:** *Efforts in the second stage of* TS *are lower, but not significantly different from efforts in the first stage. In this sense, the* TS *treatment is incentive maintaining.*

4.5 Concluding Remarks

Promotions in most hierarchical organizations take the form of multi-stage elimination tournaments. In this chapter we have studied behavior in such tournaments with simple laboratory experiments. Our results demonstrate the importance of carefully analyzing the incentive effects of promotions in multi-level hierarchies. They show that the basic logic of incentive provision in multi-stage elimination tournaments works in the sense that people take future promotion possibilities into account when deciding on current work effort. However, we also observe important departures from theoretical predictions. Subjects tend to exert excess effort in the first stage of our two-stage elimination tournament. By contrast, we do not observe this phenomenon in a strategically equivalent one-stage tournament. Under a more convex wage structure, the overprovision of effort is even more pronounced.

Our experiments suggest that behavior in multi-stage tournaments deviates from behavior in one-stage tournaments in a systematic way. Our data do, however, not allow us to give a definite answer on the precise mechanism that causes this change in behavior. Several factors may act in concert: it could be that subjects experience additional non-monetary "joy of winning" when being promoted (Parco et al. 2005), which might be more pronounced when the hierarchy has more layers. An additional potential rationale for subjects' behavior are preferences for status (Moldovanu et al. 2007). Multi-stage tournaments with their more precise definition of hierarchical

[9] The same exercise for the *TSC* treatment yields a value of 14 while the average effort level actually observed in the second stage of this treatment is 82.9. Thus, although subjects on average choose efforts below the equilibrium value of 100 in the second stage of *TSC*, this cannot explain the excess efforts exerted in the first stage.

level (and status) might trigger especially competitive behavior of status concerned agents. Our data on subjects' expectations are consistent with these interpretations: about two thirds of subjects choose an effort above the second-highest effort level that they expect from their competitors, irrespective of the absolute level of the effort expectation.

The observed behavior could also help to rationalize why firms rely so strongly on promotions as incentive device, even in work environments where more direct performance pay is feasible. Excess effort in early stages of multi-stage tournaments makes this form of incentive provision comparatively "cheap" for the principal as it decreases the wage cost per unit of effort. For instance, in our *TS* tournament this cost is 9.81 Cent in theory, but only 8.36 Cent in practice. This implies that a principal who implemented the *TS* wage scheme had to pay less for every unit of effort (and production) than theoretically predicted. Which wage profile a principal or tournament designer actually prefers depends on his objectives. In some situations it might be sufficient to concentrate on the cost per effort. In other situations a tournament designer may, for example, put special emphasis on the performance of agents in higher stages of the tournament. The *TSC* treatment suggests that excess effort in early stages eventually might come at the cost of reduced performance in later stages if the wage structure becomes too convex. Independent of the specific objective function, a tournament designer should take into account that agents' behavior can ultimately change the optimal architecture of promotion tournaments in terms of wage profiles, promotion rates, etc.

Bibliography

Abowd, J. M., Kramarz, F. and Margolis, D. N. (1999), 'High wage workers and high wage firms', *Econometrica* **67**(2), 251–333.

Adams, J. S. (1963), 'Wage inequities, productivity and work quality', *Industrial Relations* **3**, 9–16.

Adams, J. S. (1965), Inequity in social exchange, *in* L. Berkowitz, ed., 'Advances in Experimental Social Psychology', Vol. 2, Academic Press, New York, pp. 267–299.

Akerlof, G. A. (1982), 'Labor contracts as partial gift exchange', *The Quarterly Journal of Economics* **97**, 543–569.

Akerlof, G. A. and Yellen, J. L. (1990), 'The fair wage-effort hypothesis and unemployment', *The Quarterly Journal of Economics* **105**, 255–283.

Altmann, S., Dohmen, T. and Wibral, M. (2008), 'Do the reciprocal trust less?', *Economics Letters* **99**(3), 454–457.

Andrews, I. R. (1967), 'Wage inequity and job perfomance: An experimental study', *Journal of Applied Psychology* **51**, 39–45.

Ashraf, N., Bohnet, I. and Piankov, N. (2006), 'Decomposing trust and trustworthiness', *Experimental Economics* **9**, 193–208.

Baker, G., Gibbons, R. and Murphy, K. J. (1992), 'Relational contracts and the theory of the firm', *Quarterly Journal of Economics* **117**(1), 39–84.

Baker, G., Gibbs, M. and Holmström, B. (1994), 'The internal economics of the firm: Evidence from personnel data', *The Quarterly Journal of Economics* **109**(4), 881–919.

Baker, G. P., Jensen, M. C. and Murphy, K. J. (1988), 'Compensation and incentives: Practice vs. theory', *Journal of Finance* **43**(3), 593–616.

Bandiera, O., Barankay, I. and Rasul, I. (2005), 'Social preferences and the response to incentives: Evidence from personnel data', *The Quarterly Journal of Economics* **120**(3), 917–962.

Bellemare, C. and Kröger, S. (2007), 'On representative social capital', *European Economic Review* **51**, 183–202.

Belzil, C. and Bognanno, M. (2008), 'Promotions, demotions, halo effects, and the earnings dynamics of american executives', *Journal of Labor Economics* **26**(2), 287–310.

Berg, J. E., Dickhaut, J. and McCabe, K. (1995), 'Trust, reciprocity, and social history', *Games and Economic Behavior* **10**, 122–142.

Bewley, T. (1999), *Why Wages Don't Fall During a Recession*, Harvard University Press, Cambridge, MA.

Blanchflower, D. G., Oswald, A. J. and Sanfrey, P. (1996), 'Wages, profits, and rent-sharing', *Quarterly Journal of Economics* **111**(1), 227–251.

Blount, S. (1995), 'When social outcomes aren't fair: The effect of causal attribution on preferences', *Organizational Behavior and Human Decision Processes* **63**, 131–144.

Bognanno, M. L. (2001), 'Corporate tournaments', *Journal of Labor Economics* **19**(2), 290–315.

Bohnet, I. and Zeckhauser, R. (2004), 'Trust, risk and betrayal', *Journal of Economic Behavior and Organization* **55**(4), 467–484.

Bolton, G. E. and Ockenfels, A. (2000), 'Erc: A theory of equity, reciprocity, and competition', *American Economic Review* **90**(1), 166–193.

Boot, A. W. A. (2000), 'Relationship banking: What do we know?', *Journal of Financial Intermediation* **9**, 7–25.

Bowles, S. and Park, Y. (2005), ‘Emulation, inequality, and work hours: Was thorsten veblen right?’, *The Economic Journal* **115**, F397–F412.

Brown, M., Falk, A. and Fehr, E. (2004), ‘Relational contracts and the nature of market interactions’, *Econometrica* **72**(3), 747–780.

Brown, M., Falk, A. and Fehr, E. (2008), ‘Competition and relational contracts: The role of unemployment as a disciplinary device’, *IZA Discussion Paper* **3345**.

Buchan, N. R., Croson, R. T. and Solnick, S. (2008), ‘Trust and gender: An examination of behavior and beliefs in the investment game’, *Journal of Economic Behavior and Organization* **68**, 466–476.

Bull, C., Schotter, A. and Weigelt, K. (1987), ‘Tournaments and piece rates: An experimental study’, *Journal of Political Economy* **95**(1), 1–33.

Bulow, J. I. and Summers, L. H. (1986), ‘A theory of dual labor markets with application to industrial policy, discrimination, and keynesian unemployment’, *Journal of Labor Economics* **4**(3), 376–414.

Burks, S. V., Carpenter, J. P. and Verhoogen, E. (2003), ‘Playing both roles in the trust game’, *Journal of Economic Behavior and Organization* **51**, 195–216.

Burks, S. V., Carpenter, J. P. and Verhoogen, E. (2007), ‘Fairness and freight-handlers: Local labor market conditions and wage fairness perceptions in a trucking firm.’, *Industrial and Labor Relations Review* **60**, 477–498.

Campbell III, C. M. and Kamlani, K. S. (1997), ‘The reasons for wage rigidity: Evidence from a survey of firms’, *Quarterly Journal of Economics* **112**(3), 759–789.

Charness, G. (2004), ‘Attribution and reciprocity in an experimental labor market’, *Journal of Labor Economics* **22**(3), 665–688.

Charness, G. and Kuhn, P. (2007), ‘Does pay inequality affect worker effort? experimental evidence’, *Journal of Labor Economics* **25**, 693–723.

Clark, A. E. and Oswald, A. J. (1996), 'Satisfaction and comparison income', *Journal of Public Economics* **61**(3), 359–381.

Clark, A., Masclet, D. and Villeval, M.-C. (2006), Effort and comparison income: Survey and experimental evidence. GATE Working Paper 06-01.

Cohn, A., Fehr, E. and Götte, L. (2008), 'Fairness and effort–evidence from a field experiment', *mimeo, University of Zurich* .

Cox, J. C. (2004), 'How to identify trust and reciprocity', *Games and Economic Behavior* **46**(2), 260–281.

Deci, E. L. (1971), 'Effects of externally mediated rewards on intrinsic motivation', *Journal of Personality and Social Psychology* **18**(1), 105–115.

Dohmen, T. and Falk, A. (2006), 'Performance pay and multi-dimensional sorting: Productivity, preferences and gender', *IZA Discussion Paper* **2001**.

Dohmen, T., Falk, A., Huffman, D. and Sunde, U. (2006), 'The intergenerational transmission of risk and trust attitudes', *IZA Discussion Paper* **2380**.

Dufwenberg, M. and Gneezy, U. (2000), 'Measuring beliefs in an experimental lost wallet game', *Games and Economic Behavior* **30**, 163–182.

Dufwenberg, M. and Kirchsteiger, G. (2004), 'A theory of sequential reciprocity', *Games and Economic Behavior* **47**, 268–298.

Easterlin, R. A. (2001), 'Income and happiness: Towards a unified theory', *The Economic Journal* **111**, 465–484.

Ehrenberg, R. G. and Bognanno, M. L. (1990), 'Do tournaments have incentive effects?', *Journal of Political Economy* **98**(6), 1307–1324.

Erev, I., Bornstein, G. and Galili, R. (1993), 'Constructive intergroup competition as a solution to the free rider problem: A field experiment', *Journal of Experimental Social Psychology* **29**, 463–478.

Eriksson, T. (1999), 'Executive compensation and tournament theory: Empirical tests on danish data', *Journal of Labor Economics* **17**(2), 262–280.

Eriksson, T., Teyssier, S. and Villeval, M.-C. (2006), 'Self-selection and the efficiency of tournaments', *IZA Discussion Paper* **1983**.

Falk, A., Fehr, E. and Fischbacher, U. (2005), 'Driving forces behind informal sanctions', *Econometrica* **73**(6), 2017–2030.

Falk, A., Fehr, E. and Huffman, D. (2008), 'The power and limits of tournament incentives', *Working Paper, University of Bonn* .

Falk, A., Fehr, E. and Zehnder, C. (2006), 'Fairness perceptions and reservations wages–the behavioral effects of minimum wage laws', *Quarterly Journal of Economics* **121**(4), 1347–1381.

Falk, A. and Fischbacher, U. (2006), 'A theory of reciprocity', *Games and Economic Behavior* **54**(2), 293–315.

Falk, A., Huffman, D. and MacLeod, W. B. (2008), 'Institutions and contract enforcement', *IZA Discussion Paper* **3435**.

Falk, A. and Kosfeld, M. (2006), 'The hidden cost of control', *American Economic Review* **96**(5), 1611–1630.

Falk, A. and Zehnder, C. (2007), 'Discrimination and in-group favoritism in a city-wide trust experiment', *IZA Discussion Paper* **2765**.

Fama, E. F. (1980), 'Agency problems and the theory of the firm', *Journal of Political Economy* **88**(2), 288–307.

Fehr, E. and Falk, A. (1999), 'Wage rigidity in a competitive incomplete contract market', *Journal of Political Economy* **107**(1), 106–134.

Fehr, E. and Falk, A. (2002), 'Psychological foundations of incentives', *European Economic Review* **46**, 687–724.

Fehr, E. and Gächter, S. (2000), 'Cooperation and punishment in public goods experiments', *American Economic Review* **90**, 980–994.

Fehr, E., Gächter, S. and Kirchsteiger, G. (1997), 'Reciprocity as a contract enforcement device: Experimental evidence', *Econometrica* **65**, 833–860.

Fehr, E., Kirchsteiger, G. and Riedl, A. (1993), 'Does fairness prevent market clearing? an experimental investigation', *Quarterly Journal of Economics* **108**, 437–460.

Fehr, E., Kirchsteiger, G. and Riedl, A. (1996), 'Involuntary unemployment and non-compensating wage differentials in an experimental labour market', *Economic Journal* **106**, 106–121.

Fehr, E., Klein, A. and Schmidt, K. M. (2007), 'Fairness and contract design', *Econometrica* **75**(1), 121–154.

Fehr, E. and Schmidt, K. M. (1999), 'A theory of fairness, competition and cooperation', *Quarterly Journal of Economics* **114**, 817–868.

Fehr, E. and Schmidt, K. M. (2006), The economics of fairness, reciprocity and altruism – experimental evidence and new theories, *in* S.-C. Kolm and J. M. Ythier, eds, 'Handbook of the Economics of Giving, Altruism and Reciprocity', Elsevier.

Fischbacher, U. (2007), 'z-tree - zurich toolbox for ready-made economic experiments', *Experimental Economics* **10**, 171–178.

Fließbach, K., Weber, B., Trautner, P., Dohmen, T., Sunde, U., Elger, C. and Falk, A. (2007), 'Social comparison affects reward-related brain activity in the human ventral striatum', *Science* **318**, 1305–1308.

Frank, R. H. (1999), *Luxury Fever – Money and Happiness in an Era of Excess*, 1 edn, Princeton University Press, Princeton, NJ.

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

Gneezy, U. and Rustichini, A. (2000), 'Pay enough or don't pay at all', *Quarterly Journal of Economics* **115**, 791–810.

Goux, D. and Maurin, E. (1999), ‘Persistence of interindustry wage differentials: A reexamination using matched worker-firm panel data’, *Journal of Labor Economics* **17**(3), 492–533.

Greiner, B. (2003), An online recruitment system for economic experiments, *in* K. Kremer and V. Macho, eds, ‘Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63’, Ges. für Wiss. Datenverarbeitung, Göttingen, pp. 79–93.

Hannan, R. L., Kagel, J. H. and Moser, D. V. (2002), ‘Partial gift exchange in an experimental labor market: Impact of subject population differences, productivity differences, and effort requests on behavior’, *Journal of Labor Economics* **20**(4), 923–951.

Harbring, C. and Irlenbusch, B. (2003), ‘An experimental study on tournament design’, *Labour Economics* **10**, 443–464.

Harbring, C. and Irlenbusch, B. (2008), ‘How many winners are good to have? on tournaments with sabotage’, *Journal of Economic Behavior and Organization* **65**, 682–702.

Hillman, A. L. and Riley, J. G. (1989), ‘Politically contestable rents and transfers’, *Economics and Politics* **1**(1), 17–39.

Holmström, B. (1982), ‘Moral hazard in teams’, *Bell Journal Of Economics* **13**(2), 324–340.

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

Holmström, B. and Milgrom, P. (1991), ‘Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design’, *Journal of Law, Economics, and Organization* **7**, 24–52.

Homans, G. C. (1961), *Social Behavior: Its Elementary Forms*, Routledge & Kegan Paul, London.

Kelley, H. H. and Stahelski, A. J. (1970), 'Social interaction basis of cooperators' and competitors' beliefs about others', *Journal of Personality and Social Psychology* **16**, 66–91.

Knack, S. and Keefer, P. (1997), 'Does social capital have an economic payoff? a cross-country investigation', *Quarterly Journal of Economics* **112**(4), 1251–1288.

Knez, M. and Simester, D. (2001), 'Firm-wide incentives and mutual monitoring at continental airlines', *Journal of Labor Economics* **19**(4), 743–772.

Knoeber, C. R. and Thurman, W. N. (1994), 'Testing the theory of tournaments: An empirical analysis of broiler production', *Journal of Labor Economics* **12**(2), 155–179.

Konow, J. (2003), 'Which is the fairest one of all? a positive analysis of justice theories', *Journal of Economic Literature* **41**, 1188–1239.

Krueger, A. B. and Summers, L. H. (1988), 'Efficiency wages and the inter-industry wage structure', *Econometrica* **56**(2), 259–293.

Köszegi, B. and Rabin, M. (2006), 'A model of reference-dependent preferences', *Quarterly Journal of Economics* **121**(4), 1133–1165.

Lawler, E. E. (1967), 'Effects of hourly overpayment on productivity and work quality', *Journal of Personality and Social Psychology* **10**, 306–313.

Layard, R. (2005), *Happiness: Lessons from a New Science*, 1 edn, Penguin, London, UK.

Lazear, E. P. (1989), 'Pay equality and industrial politics', *Journal of Political Economy* **97**(3), 561–80.

Lazear, E. P. (2000), 'Performance pay and productivity', *American Economic Review* **90**(5), 1346–1361.

Lazear, E. P. and Gibbs, M. (2008), *Personnel Economics in Practice*, 2 edn, John Wiley & Sons.

Lazear, E. P. and Rosen, S. (1981), 'Rank-order tournaments as optimum labor contracts', *Journal of Political Economy* **89**(5), 841–864.

Loewenstein, G., Thompson, L. and Bazerman, M. (1989), 'Social utility and decision making in interpersonal contexts', *Journal of Personality and Social Psychology* **57**, 426–441.

MacLeod, W. B. and Malcomson, J. M. (1989), 'Implicit contracts, incentive compatibility, and involuntary unemployment', *Econometrica* **57**(2), 447–480.

MacLeod, W. B. and Malcomson, J. M. (1998), 'Motivation and markets', *American Economic Review* **88**(3), 388–411.

Maximiano, S., Sloof, R. and Sonnemans, J. (2007), 'Gift exchange in a multi-worker firm', *The Economic Journal* **117**, 1025–1050.

Medoff, J. L. and Abraham, K. G. (1980), 'Experience, performance, and earnings', *The Quarterly Journal of Economics* **95**(4), 703–36.

Milgrom, P. and Roberts, J. (1992), *Economics, Organization and Management*, Prentice Hall, Upper Saddle River, New Jersey.

Moldovanu, B. and Sela, A. (2006), 'Contest architecture', *Journal of Economic Theory* **126**(1), 70–96.

Moldovanu, B., Sela, A. and Shi, X. (2007), 'Contests for status', *Journal of Political Economy* **115**(2), 338–363.

Mowday, R. T. (1991), Equity theory predictions of behavior in organizations, *in* R. M. Steers and L. W. Porter, eds, 'Motivation and Work Behavior', McGraw-Hill, New York, pp. 111–131.

Neumark, D. and Postlewaite, A. (1998), 'Relative income concerns and the rise in married women's employment', *Journal of Public Economics* **70**, 157–183.

O'Donoghue, T. and Rabin, M. (1999), 'Incentives for procrastinators', *Quarterly Journal of Economics* **114**, 769–816.

Orrison, A., Schotter, A. and Weigelt, K. (2004), 'Multiperson tournaments: An experimental investigation', *Management Science* **50**(2), 268–279.

Parco, J. E., Rapoport, A. and Amaldoss, W. (2005), 'Two-stage contests with budget constraints: an experimental study', *Journal of Mathematical Psychology* **49**, 320–338.

Pfeffer, J. and Langton, N. (1993), 'The effect of wage dispersion on satisfaction, productivity, and working collaboratively: Evidence from college and university faculty', *Administrative Science Quarterly* **38**(3), 382–407.

Prendergast, C. (1999), 'The provision of incentives in firms', *Journal of Economic Perspectives* **37**, 7–63.

Reuter, M., Montag, C., Altmann, S., Bendlow, F., Elger, C., Kirsch, P., Becker, A., Schoch, S., Simon, M., Weber, B. and Falk, A. (2008), 'Genetically determined differences in human trust behavior: The role of the oxytocin receptor gene', *mimeo, University of Bonn* .

Rosen, S. (1986), 'Prizes and incentives in elimination tournaments', *American Economic Review* **76**(4), 701–715.

Saint-Paul, G. (1996), *Dual Labor Markets – A Macroeconomic Perspective*, 1 edn, MIT Press, Cambridge, MA.

Schotter, A. and Weigelt, K. (1992), 'Asymmetric tournaments, equal opportunity laws, and affirmative action: Some experimental results', *The Quarterly Journal of Economics* **107**(2), 511–539.

Shapiro, C. and Stiglitz, J. E. (1984), 'Equilibrium unemployment as a worker discipline device', *American Economic Review* **74**(3), 433–444.

Taylor, C. R. (1995), 'Digging for golden carrots: An analysis of research tournaments', *American Economic Review* **85**(4), 872–890.

Thöni, C. and Gächter, S. (2008), ''Kinked conformism' in voluntary cooperation', *mimeo* .

Torgler, B., Schmidt, S. L. and Frey, B. S. (2006), 'Relative income position and performance: an empirical panel analysis', *IEW Working Paper* **268**.

Appendices

A.1 Instructions for Chapter 1

In this part of the experiment, your endowment, contributions, and earnings are calculated in points. At the end of the experiment the points you have earned will be converted at an exchange rate of **50 points = 1 Euro** and paid out to you in cash.

In this part of the experiment you are in a group together with one of the other 15 participants. This participant will be **assigned to you randomly by the computer**. In each group of two there are two roles (a sender and a receiver). One member of the group will be the sender, and one member of the group will be the receiver. Each participant receives an endowment of 120 points.

The experiment consists of two stages:

In the first stage the sender can make a transfer to the receiver. The transfer is a number between 0 and 120 in steps of 20 points, i.e., either 0, 20, 40, 60, 80, 100, or 120 points. The transfer is tripled by the experimenter.

Examples: In case the sender transfers 60 points, the receiver gets 180 points. In case the sender transfers 10 points, the receiver gets 30 points. In case nothing is transferred, the receiver gets 0 points etc.

This means that, at the end of the first stage, the amount of points available to the receiver is the sum of the endowment plus three times the transfer.

In the second stage the receiver can transfer any number of points back to the sender. The amount sent back is not tripled. The back transfer must be a number between 0 and 480. Once the receiver has decided on the back transfer the payoffs are determined.

The payoff functions are thus as follows:

For the sender: 120 – transfer + back transfer

For the receiver: 120 + 3*transfer – back transfer

An example: Let's assume the sender transfers 40 points. At the end of the first stage the sender then has 120 – 40 = 80 points, and the receiver has 120 + 3*40 = 240 points. In the second stage, the receiver chooses a back transfer of 50 points. The payoffs then are as follows: for the sender: 120 – 40 + 50 = 130 points. For the receiver: 120 + 3*40 – 50 = 190 points.

At the end of this part of the experiment, the computer randomly determines if you are a sender or a receiver. Because you do not know if you are a sender or a receiver by now, you have to make a decision for both roles. Therefore you have to make decisions on two separate computer screens: at first for the role of the receiver, then for the role of the sender.

Please press CONTINUE to make your decision.

B.2 Instructions for Chapter 2 (EWT)

Welcome to this decision-making experiment.

Please read the following instructions carefully. At the end of the instructions you will find several control questions. After correctly answering these questions the experiment will begin. The experiment will be conducted anonymously, that is to say you will not learn with whom of the other participants you are interacting. During the experiment you can earn taler. How much you earn depends on your decisions and the decisions of the other participants. At the end of the experiment these taler will be **converted to Euro at an exchange rate of**

1 Taler = 1 Cent

and be paid out to you. Please wait until you will be called to collect your earnings. Please turn in all the documents you received to conduct the experiment when you collect your earnings.

You will begin the experiment with an initial endowment of 400 talers (4 Euro). This amount of talers will increase or decrease respectively with your earnings from the different rounds. You can always rule out losses through your own decisions.

Please keep in mind that from now on and throughout the whole experiment you are not allowed to talk to the other participants. In case you do not respect this rule we are forced to terminate the experiment. **If you have any questions, please hold your hand out of the cubicle and we will come to you.**

As a participant in this experiment you either play the role of an **employer** or an **employee**. Your specific role will be randomly determined at the beginning of the experiment and it will be the same for all the following rounds.

The experiment will last for 12 rounds. At the beginning of each round, groups of 3 participants will be formed, each consisting of one employer and two employees (employee1 and employee2). **In every new round new groups of three participants will be randomly formed.** Your decisions will only be disclosed to the other two participants in your current group. All the other participants will not be informed about your decisions.

Each round consists of two stages. **In the first stage employee1 and employee2 independently choose their amounts of work.** There is a set of ten different possible amounts of work out of which they can select their amount of work. **The smallest amount of work is 1 and the highest is 10**. Every unit of work creates 10 taler for the employer. For example: If the selected amout of work is 1 then the employer receives 10 taler, if it is 2, he will receive 20 taler and so on. If the selected amount of work is 10 the employer will receive 100 taler. For the employee the selection of an amount of work is associated with costs. The higher the selected amount of work, the higher the associated costs for the employee. The employee's costs only depend on his selected amount of work. The amount of work selected by the other employees does not influence his costs. The costs regarding the amount of work for the employee are the following:

Amount of work e_i	1	2	3	4	5	6	7	8	9	10
Cost $c(e_i)$	0	1	2	4	6	8	10	13	16	20

So, the amount 1 of work is free for the employee. If the amount of work is 2, it will cost the employee 1 taler and so on. If the amount of work is 10 the associated cost is 20. All employees have the same cost table. The cost table will not change during the 12 rounds.

In the second stage the employer will be informed about the amounts of work selected by the employees in his group. Subsequently the **employer will chose a wage w**, which he pays out to employee1 and employee2, that is he pays w taler to employee1 and w taler to employee2. The wage per capita must not be smaller than 0 taler and must not be greater than 100 taler.

While the employer chooses the wage-level w, each employee estimates the amount of work **the other employee** in his group has selected. If his estimation proofs correct he will receive 2 taler, if incorrect 0 taler. Other than that this estimation has **no further consequences**. This estimation will not be disclosed to any of the other participants.

At the final stage of the round both employees and the employer will be informed about the selected amounts of work, the payroll of each employee and the resulting earnings of the employer.

The earnings in taler per round are thus calculated in the following way:

Earnings of the employer = 10 x amount of work of employee1 + 10 x amount of work of employee2 - 2 x wage (w)

Earnings of employee 1 = wage (w) - cost of the amount of work of employee1

Earnings of employee 2 = wage (w) - cost of the amount of work of employee2

At the end of the experiment your total earnings, comprising the initial endowment, the sum of your earnings in the different rounds and your earnings for correctly estimating the amounts of work will be paid out to you. As mentioned above 1 taler equals 1 Cent.

Please note that you are to write down all data of your group (amount of work, wage w, earnings) on the result sheet at the end of each round.

C.3 Predictions for Fair-Minded Players

In this appendix, we analyze whether it can be profitable for firms to offer fair, rent-sharing contracts in the IC treatment when we assume that a fraction of workers is fair-minded. Our aim is not to fully characterize the set of feasible (Bayesian) Nash Equilibria. Rather, we want to illustrate how fairness preferences shape the profitability of certain contracting strategies and how this could influence labor market outcomes. In particular, we aim at illustrating how contractual incompleteness can render job rationing optimal for firms and, thus, can give rise to unemployment.

We first analyze the case where fair-minded workers insist on an equal distribution of surplus among all members of a firm. It is shown that under these preferences, it can be beneficial for firms to build up long-term work relationships and offer strictly positive rents throughout the game. However, in this situation it is always more profitable for firms to ration jobs and employ only one worker. We then analyze the case where a fair-minded worker requires an equal share of the surplus generated by himself. It is shown that—under this weaker fairness benchmark—firms always prefer to employ the maximum number of workers.

C.3.1 Strictly Egalitarian Fairness Preferences

We assume that a fraction $0 \leq p \leq 1$ of workers has "strictly egalitarian" fairness preferences. Workers' types are private information. A worker with such preferences is assumed to fulfill the contract offered to him as long as the offered wage and desired level of effort are such that the worker receives at least an equal share of the total (net) surplus generated if all workers employed by the firm receive the same contract offer and fulfill their contract. If the contract terms yield lower material payoffs to the worker, the worker will accept the contract offer and shirk by providing the minimum possible effort.[10] The utility function of a worker with such preferences

[10]Since rejecting a contract offer yields a material payoff of zero and since the cost of providing the minimum effort is also zero, we can assume without loss of generality that all contract offers where $w > 0$ are accepted.

can be summarized as follows:

$$u(w, e, \hat{e}) = \begin{cases} w - c(e) & \text{if} \quad w - c(\hat{e}) < \frac{1}{n+1} n[q(n)\hat{e} - c(\hat{e})] \\ w - c(e) - b\max[\hat{e} - e; 0] & \text{if} \quad w - c(\hat{e}) \geq \frac{1}{n+1} n[q(n)\hat{e} - c(\hat{e})] \end{cases}$$

n denotes the number of workers employed by the worker's firm, $q(n)$ is the productivity of effort of a worker employed in a n-worker firm, $\hat{e}$ and e are the contractually desired and actually exerted levels of effort, $c(\hat{e})$ and $c(e)$ are the effort cost of the desired and exerted level of effort, and w is the wage stipulated in the contract. The assumption regarding the fairness parameter b are as follows:

$$b > max\{c'(e)\}$$

This implies that a fair-minded worker maximizes utility by fulfilling the contract (i.e., choosing $e = \hat{e}$) as long as he gets (at least) an equal share of the net surplus that is created if all n workers of the firm fulfill their contract. In case the contract offers him less than a fraction $\frac{1}{n+1}$ of net surplus, the worker will always choose the minimum possible level of effort.

In what follows, we will show that for these fairness preferences there can be no equilibrium where (i) all firms employ two workers and (ii) efforts exceed the minimum possible level. To do so, we first analyze under which conditions two-worker firms would be willing to pay strictly positive rents to workers in the final period of the game. Next we calculate firms' and workers' payoffs in pre-final periods, analyzing the requirements to sustain cooperation in these periods. We then analyze the case of one-worker firms. Finally, we show that for any fraction of p, firms are better off if they employ only one worker.

Two-Worker Firms

Final Period

Firms are assumed to maximize (expected) monetary profits. In the final period, T, a firm has to decide between offering payoff-equalizing contracts to workers (and receiving above-minimal effort only from the fraction of fair workers) and the outside option of offering an uncooperative contract where both types of workers shirk.

Since unemployment benefits for workers are equal to 0, workers will accept any contract that offers them at least a wage of 1. Therefore, the outside option of a firm is to offer a wage of 1 and desire work effort of 1. This outside option ensures the following last period profits Π_T^O for firms:

$$\Pi_T^O(w, \hat{e}) = \Pi_T^O(1, 1) = nq(n) * 1 - n1$$

It is easy to see that—with the parameters specified in Section 3.2.3—firms who offer the uncooperative contract $[1, 1]$ are always better of if they employ $n = 2$ workers. Since $q(2) = 7$, a firm hiring two workers at $[w, \hat{e}] = [1, 1]$ earns $\Pi_T^O = 12$ whereas a one-worker firm earns only $\Pi_T^O = 9$ (since $q(1) = 10$).

Can a firm do better than this in period T by hiring two workers at a fair, i.e., payoff-equalizing wage? Since all fair workers fulfill such a contract in T and all selfish workers shirk, the expected profits of a firm employing two workers at payoff equalizing terms $[w, \hat{e}]$ look as follows:

$$\Pi_T(w, \hat{e}) = p^2 2[q(2)e(w) - w] + (1-p)^2 2(q(2)*1 - w) + p(1-p)2[q(2)e(w) + q(2)*1 - 2w]$$

The first term specifies the firm's profit if the firm happens to employ two fair-minded workers (which occurs with probability p^2), the second term specifies the profits in case both workers are selfish (occurring with probability $(1-p)^2$), and the final term summarizes the cases where one worker is fair-minded and fulfills the contract and one worker is selfish (which occurs with probability $p(1-p)2$).

To make offering a strictly positive rent profitable for firms, it is necessary that $\Pi_T(w, \hat{e}) > \Pi_T^O = 12$. From this, we can derive a threshold value p^* for the fraction of fair workers above which it is profitable for firms to offer payoff equalizing contracts in period T. Under our assumptions regarding the production function, effort cost, and workers' utility, this threshold value is $p^* \approx 0.82$. If $p < p^*$ no firm will offer other contract terms than $[w, \hat{e}] = [1, 1]$ in period T. If $p \geq p^*$, two-worker firms will offer the "fair" contract $[w, \hat{e}] = [53, 10]$. Any intermediate wage-effort combination will never be offered by a two-worker firm in period T, i.e., it is not optimal to desire $1 < \hat{e} < 10$ in period T at a payoff-equalizing wage.

Pre-Final Periods

Analyzing optimal contracts in pre-final periods is relatively straightforward. A fair worker will cooperate in any period t as long as the contract offered is at least sharing surplus equally. A selfish worker will cooperate and fulfill a contract in pre-final periods if expected *future* rents give him higher payoffs than the immediate gains from shirking. If $p > p^*$, firms will offer the fully cooperative contract $[w, \hat{e}] = [53, 10]$ in period T. If a firm re-employs its workers in period $t+1$ as long as they fulfill the stipulated contract in period t and dismisses all workers who do not do so, a selfish worker in period $T-1$ has to decide between shirking in period $T-1$ and being unemployed in period T vs. adhering to the contract in $T-1$ and being re-employed in period T. The first option yields a payoff of $\pi_{T-1} + \pi_T = w_{T-1} - c(1) + 0$, the latter option yields $\pi_{T-1} + \pi_T = w_{T-1} - c(\hat{e}_{T-1}) + 53$ since a re-employed selfish worker receives $w_T = 53$ in period T and shirks (i.e., $e_T = 1$). As $c(\hat{e}_{T-1}) < 53$, selfish workers who are offered $[w, \hat{e}] = [53, 10]$ in period $T-1$ will fulfill their contract.

Therefore it is optimal for two-worker firms to offer $[w, \hat{e}] = [53, 10]$ in all pre-final periods $t < T$ if $p > p^*$. Under this contract, both types of workers fulfill their contracts in $t < T$ and earn $\pi_t = 53 - c(10) = 35$ and firms earn $\Pi_t = 2 * q(2) * 10 - 2 * 53 = 34$. In the final period, fair workers fulfill and earn $\pi_T = 35$, selfish workers shirk and earn $\pi_T = 53$, and firms earn $\Pi_T = p^2 * (34) + (1-p)^2 * (-84) + p(1-p) * 2 * (-27) > \Pi_T^O$.

One-Worker Firms

Final Period

Can it be profitable for a firm to employ only one worker in period T and offer him a strictly positive rent? To analyze this question, remember that a firm's outside option is hiring two workers at the minimum possible wage of 1. If a firm instead hires one worker at a payoff equalizing wage level, its expected final-period profits are given by

$$\Pi_T(w, \hat{e}) = p[q(1)e(w) - w] + (1-p)[q(1) * 1 - w]$$

We can again calculate a threshold share of fair workers, p^{**}, above which it is profitable for a firm to offer a fair, cooperative contract to one worker in the final period. Evaluating Π_T and Π_T^O for our parameter constellation yields the threshold $p^{**} \approx 0.68$. Above p^{**}, a firm employing one worker and offering the payoff-sharing contract $[w, \hat{e}] = [59, 10]$ receives higher expected profits than Π_T^O. Again, it is not profitable for firms to offer fair contracts at intermediate desired effort levels $1 < \hat{e} < 10$.

Pre-Final Periods

Turning to pre-final periods and applying the results for two-worker firms, cooperation by selfish and fair-minded workers can be sustained in every period $t < T$ if $p > p^{**}$. If this is the case, a firm can offer the fully cooperative rent-sharing contract $[w, \hat{e}] = [59, 10]$ to one worker in every period and apply a strategy of contingent contract renewal, i.e., re-employing its worker if he adhered to the contract in the period before and dismissing the worker if $e \neq \hat{e}$ in the previous period. Fair workers adhere to this contract in every period since it offers them a fair share of the surplus. Selfish workers imitate fair workers in all pre-final periods since the prospect of being re-employed and earning future rents gives them an incentive to do so in all but the final period. In the final period, selfish workers provide minimum effort and earn $\pi_T = w = 59$.

Comparing One-Worker Firms and Two-Worker Firms

Comparing the results for the one-worker and two-worker case, it is immediately obvious that there can be no Bayesian Nash Equilibrium with above-minimal cooperation and full employment (i.e., all firms employ two workers). The threshold share of fair workers necessary for inducing cooperation in the final period (and consequently in all pre-final periods) is higher for the two-worker firm than for a one-worker firm, i.e., $p^* > p^{**}$. If the actual level of fair minded workers is between the two thresholds, cooperation can only be achieved if just one worker is employed by the firm.

Moreover, even if there are sufficiently many fair-minded workers to make cooperation in two-worker firms feasible (i.e., if $p \geq p^*$), firms are better of if they employ only one worker. The fair contract which induces maximum effort in the two-worker firm is $[w, \hat{e}] = [53, 10]$. This contract yields firm profits of $\Pi_t = 34$ in all periods $t < T$ and $\Pi_T(n = 2) \geq \Pi_T^O$ in T. By contrast, the fair and fully efficient contract in a one-worker firm is $[w, \hat{e}] = [59, 10]$. This contract generates strictly higher firm profits in every period: $\Pi_t = 41$ in $t < T$ and $\Pi_T(n = 1) \geq \Pi_T(n = 2)$ for all values of p.

There are several reasons for why cooperation is harder to achieve in two-worker firms. First, two-worker firms face a higher probability of meeting *at least one* selfish worker who definitely shirks in final period: $1 - p^2 > 1 - p$. Moreover, surplus to be distributed is not twice as high as in one-worker firms due to decreasing returns to scale ($q'(n) < 0$). Both effects increase the threshold level of fair-minded workers necessary to induce cooperation in the final period for two-worker firms. The latter effect in addition decreases profits of two-worker firms in every period compared to the profits of a one-worker firm. For our assumptions on the production function, receiving 33% of the higher total (net) surplus generated in a two-worker firm is less profitable than receiving 50% of the (smaller) surplus produced in a one-worker firm.

C.3.2 Relation-Specific Egalitarian Fairness Preferences

How do outcomes change if workers' fairness preferences are not characterized by the strict egalitarianism stipulated above? In what follows, we stick to the assumption that a fraction p of workers is fair-minded. However, we now assume that a fair-minded worker considers a contract offer fair if it splits the surplus *generated by himself* equally between him and the firm. The fair-minded worker's utility function is therefore given by:

$$u(w, e, \hat{e}) = \begin{cases} w - c(e) & \text{if} \quad w - c(\hat{e}) < q(n)\hat{e} - c(\hat{e}) \\ w - c(e) - b\max[\hat{e} - e; 0] & \text{if} \quad w - c(\hat{e}) \geq q(n)\hat{e} - c(\hat{e}) \end{cases}$$

If the firm employs only one worker, this notion of fairness is equivalent to the one assumed above. However, firms employing two workers can now extract up to 50% of the net surplus generated by workers' efforts, compared to only 33% in the strictly egalitarian case. If other words, a fair-minded worker neglects his firm's payoff from relations with other workers, and considers only the distribution of payoffs in his own work relationship.

As a consequence, the threshold fraction of fair-minded workers necessary to give two-worker firms an incentive to offer positive rents in the final period of the game decreases. Following the steps of the calculation above, it now pays off for firms to employ two workers in the final period and offer them fair contract terms as soon as $p > p^{***} \approx 0.68$. At this critical value, a two-worker firm's best strategy in the final period is to offer the fair contract $[w, \hat{e}] = [34, 8]$. In pre-final periods, the optimal strategy for a two-worker firm is to offer the fully cooperative contract $[w, \hat{e}] = [44, 10]$.

Since the "relation-specific" fairness benchmark implies that fair-minded workers in one-worker firms still require 50% of net surplus, the critical value for profitably employing one worker and the optimal contracts for one-worker firms remain unchanged from the previous analysis. Importantly, this implies that under the weaker fairness benchmark it is *always* better for firms to employ *two* workers for our experimental parameters. If $p < p^{***}$, both types of firms can only implement minimal effort in all periods. Since the second worker is productive, a firm prefers employing two workers at the same effort level for the minimal wage $w = 1$. If $p > p^{***}$, a firm employing one worker receives 50% of $q(1) * e - c(e)$, whereas a firm employing two workers receives 50% of $2 * [q(2) * e - c(e)]$. In every pre-final period, the maximum effort $e = 10$ will be provided in both types of firms. This means that a two-worker firm's payoff is equal to 50% of the first-best surplus, whereas a one-worker firms payoff is equal to the same share of the lower surplus that can be achieved with one worker. Similarly, in the final period of the game, payoffs for a two-worker firm are strictly higher than the one for a one-worker firm for every level of $p > p^{***}$.

C.4 Instructions for Chapter 3 (IC Treatment)

In what follows, we present a translation of the instructions for buyers (i.e., employers) in the IC treatment. The instructions for workers in this treatment had a similar structure. The instructions of participants in the C treatment differed only in the description of the second stage (i.e., the work phase).

Instructions for Buyers

You are now taking part in an economic experiment. Please read the following instructions carefully. Everything that you need to know to participate in this experiment is explained below. Should you have any difficulties in understanding these instructions please raise your hand. We will answer your questions at your cubicle.

At the beginning of the experiment you will receive an initial endowment of **8 Euros**. Over the course of the experiment you can increase your income by earning **points**. The amount of points that you earn during the experiment depends on your decisions and the decisions of other participants.

All points that you earn over the course of the experiment will be exchanged into Euros at the end of the experiment. The exchange rate will be:

1 Point = 4 Cents

At the end of the experiment, the amount of money that you earned during the experiment as well as your 8 Euros initial endowment will be paid out in cash.

The experiment is divided into several periods. In each period you have to make decisions which you will enter in a computer. In total, there will be 18 periods.

Please note that communication between participants is strictly prohibited during the experiment. In addition we would like to point out that you may only use the computer functions which are required for the experiment. Violation of these rules will lead to exclusion from the experiment. In case you have any questions we shall be glad to assist you.

Prior to the experiment the 24 participants were divided into 2 groups: buyers and sellers. In this experiment there are 7 buyers and 17 sellers.

You will be a buyer for the entire duration of the experiment. All participants have received an identification number which they will keep for the entire experiment. Your identification number is stated on the documentation sheet in front of you.

Short Overview of the Experimental Procedures

In each period of the experiment every buyer can trade a product with no, one or two sellers. The seller earns a profit through the trade when he sells the product at a price which exceeds his production costs. The buyer earns a profit through the trade when the price he pays for the product is less than what the product is worth to him. The production costs of the traded product as well as the buyer's valuation of the product depends on the quality of the product. In addition the value of the product for the buyer depends on the number of products bought. Two products of a certain quality are worth more to the buyer, but not worth twice as much as one product of the same quality.

The experiment lasts 18 periods. In each period the procedures are as follows:

Each period commences with a **trading phase** which lasts 200 seconds. During this phase buyers can submit trade offers which can be accepted by sellers. When submitting an offer a buyer has to specify three things:

- Which price he offers to pay
- which product quality he desires
- and finally, which sellers he wants to submit the offer to. Buyers can submit two types of offers: private offers and public offers. **Private offers are submitted to one specific seller** and can only be accepted by that seller.

Public offers are submitted to all sellers and can be accepted by any seller.

As a buyer you can submit as many offers as you like in each period. Once submitted, offers can be accepted constantly. **Each seller can only enter one trade agreement in each period. Each buyer can at most enter two trade agreements.** As there are 7 buyers and 17 sellers, some sellers will not trade in each period.

After the trading phase, every sellers who accepted a trade agreement has to determine which quality of product he will provide to his buyer. **Hereby, the seller is not obliged to provide the product quality desired by the buyer.**

Once every seller has chosen which product quality to provide, earnings of all participants for the given period are determined. Subsequently, the next period commences. The earnings from all 18 periods will be summed up at the end of the experiment, exchanged into Euros and paid out in cash together with your initial endowment.

The Experimental Procedures in Detail

There are 7 buyers and 17 sellers in the experiment. You are a **buyer** for the entire duration of the experiment. During the experiment you will enter your decisions in a computer. In the following we describe in detail how you can make your decisions in each period.

1. The Trading Phase

Each period commences with a trading phase. During the trading phase the buyers can enter into trading agreements with the sellers. In order to do so, **each buyer can submit as many trade offers as he wishes.** In each trading phase you will see the following screen.

In the top left corner of the screen you will see the current period of the experiment. In the top right corner of the screen you will see the time remaining in this

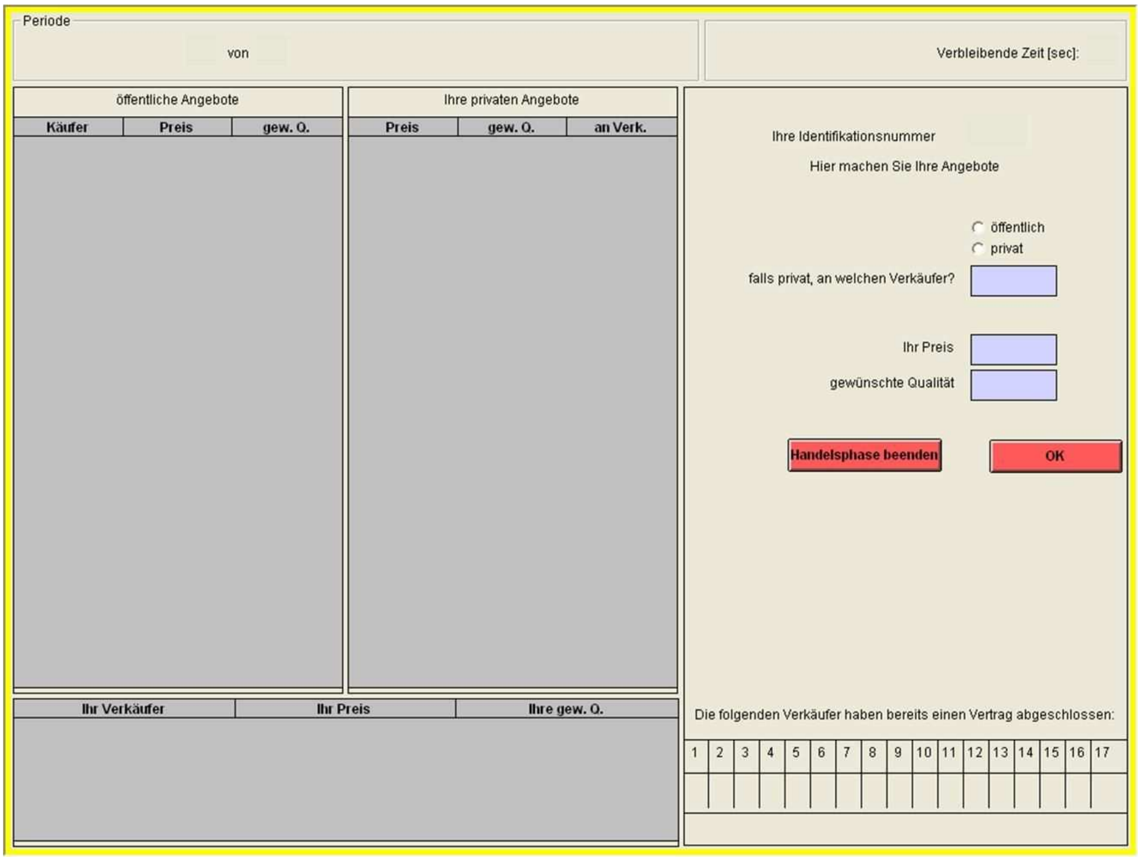

trading phase, displayed in seconds. **The trading phase in each period lasts 200 seconds.** When this time is up the trading phase is over. Hereafter, no further offers can be submitted or accepted for this period.

Once you see the above screen displayed the trading phase commences. As a buyer you now have the opportunity to submit trade offers to the sellers. In order to do so you have to enter three things on the right hand side of the screen:

(A) First, you have to specify whether you want to submit a public or private offer:

Public trade offers

Public offers will be communicated to all participants in the market. All sellers see all public offers on their screens. A public offer can therefore be accepted by **any seller**. As a buyer you will also see all public offers submitted by all buyers. If you want to submit a public offer, please mark the field "public" using the mouse.

Private trade offers

Private offers are submitted to **one seller** only. Only this seller will be informed about this offer and only this seller can accept the trade offer. No other seller or buyer will be informed about that offer. If you want to submit a private offer, please mark the field “private” using the mouse. In the field below, you then have to specify **which seller you want to submit the offer to.** Each of the 17 sellers has an identification number (Seller1, Seller2, ..., Seller17). Each seller keeps his identification number for the entire duration of the experiment. To submit an offer to a specific seller you enter the number of that seller (e.g. “4” for Seller4).

(B) Once you have specified who you want to submit an offer to, you must determine **which price you offer.** You enter the offered price into the field “your price”. The price you offer must not be below 0 or above 100:

$$0 \leq \textbf{Price offered} \leq 100$$

(C) Finally, you have to specify which product quality you desire. You enter this in the field “desired quality”. **Your desired quality** cannot be lower than 1 or higher than 10.

$$1 \leq \textbf{Desired quality} \leq 10$$

After you have completely specified your trade offer, you have to click the “ok” button to submit it. As long as you have not clicked “ok” you can still change your trade offer. After you click “ok” the offer will be displayed to all sellers you have submitted it to.

On the left side of your screen you see the heading “public offers”. All public offers in the current trading phase will be displayed here—your public offers as well as the public of all other buyers. You can see which buyer submitted the offer, which price he offered and which quality he desired. All buyers also have an identification number which they keep throughout the experiment (Buyer1, Buyer2, ..., Buyer7).

In the middle of your screen, under the header “your private offers” you will see all private offers which you have submitted in the current trading phase. Here you can see which seller you submitted an offer to, which price you offered and which quality you desired.

As long as none of your offers has been accepted by a seller, you as a buyer can submit as many private and public offers as you wish in each period. Each offer that you submit can be accepted at any time during the trading phase.

As soon as one of your offers has been accepted you are informed which seller accepted which of your offers. In the bottom left corner of your screen the identification number of the seller who accepted the offer will be displayed as well as your offered price and desired quality. At the same time all your other offers will be automatically canceled.

You can then decide if you want to enter another trade agreement. **Each buyer can enter no, one or two trade agreements in each period.** If you want to enter another trade agreement you can submit further offers to the sellers. As long as none of your offers has been accepted by a seller you can offer as many private and public offers for the second trade as you wish.

If you do not want to enter another trade agreement you can press the button "finish trading phase". This reduces the length of the trading phase if other buyer wants to submit further offers. By pressing the button, the offers you have already submitted will be automatically canceled and you can not submit any further offers. Trade agreements which were already accepted by a seller of course persist. In addition you will continue to see the screen of the trading phase until it is definitely over.

No seller can enter more than one trade agreement in each period. You will be constantly informed which sellers have not yet entered a trade agreement. In the table with the title "The following sellers have already entered a trading agreement" you can see 17 fields. Once a seller has accepted an offer a "+" will appear in the field next to his identification number. You cannot submit private offers to a seller who has already accepted an offer.

The trading phase is over after 200 seconds have elapsed, or once all buyers have entered two trade agreements, or if the remaining buyers have signalized that they do not want to enter trade agreements anymore by pressing the button "finish trading phase".

No buyer is obliged to submit trade offers, and no seller is obliged to accept a trade offer.

2. Determination of actual product quality

Following the trading phase, all sellers who have entered a trade agreement determine which product quality they will supply to their respective buyer. First, the sellers see again the price and the desired quality on a new screen. If you have entered two trade agreements in one period your sellers can also see the price and the desired quality of your other seller. The sellers then decide independently which actual product quality to choose for their product. **The product quality which you desired in your trade offer is not binding for your seller(s).** Your seller can choose exactly the quality you desired, but he can also choose a higher or lower product quality. The product quality which your seller chooses has to be an integer between 1 and 10:

$$1 \leq \textbf{Actual product quality} \leq 10$$

While the sellers determine the actual product quality, we ask you on a separate screen to specify which quality(ies) you expect him (them) to supply. In addition we ask you to state how sure you are about this expectation.

How are incomes calculated?

Your income:

If you **do not enter a trade agreement** during a trading phase you earn an income of **0 points** in this period.

If you have entered **one** trade agreement, your income depends on which price you offered and which product quality your seller supplied to you. Your income will equal 10 times the actual product quality minus the price you pay. Your income will thus be determined as follows:

Your income = 10*Actual product quality - Price

If you have entered **two** trade agreements, your income depends on which prices you offered to both sellers and which product qualities were supplied to you by the sellers. The value of the products **in total** can be higher for you if you enter two trade agreements but the value of **a single** product is lower.

In other words, two products of a certain quality are worth more to you, but not worth twice as much as one product with the same quality. If you buy **one** product you earn 10 times the chosen product quality. If you buy **two** products you earn 7 times the quality of the first product and 7 times the quality of the other product. Of course, when you buy two products you also have to pay two prices. Your income if you enter two trade agreements thus is determined as follows:

Your income = 7*Actual product quality product 1 + 7*Actual product quality product 2 - Price 1 - Price 2

An example: If you enter **one** trade agreement and the actual product quality is 8, your income is 80 minus the price. If you enter **two** trade agreements and both actual product qualities are 8 your income is 112 (=7*8+7*8) minus both prices. However, if—for instance—one actual product quality is 8 and the other quality is 1, your income is 63 (= 7*8 + 7*1) minus both prices.

As you can see from the above formula your income is higher, the higher the product quality actually supplied to you by the seller(s). At the same time your income is higher, the lower the price(s) you have to pay for the product(s).

Income of your seller:

If a seller has not entered a trade agreement during a trading phase he earns an income of 0 points in this period.

If a seller has accepted a trade offer his income will equal the price he receives minus the production costs he incurs for the product supplied. The income of your seller is determined as follows:

Income of your seller = Price - production cost

The production costs of a seller are higher, the higher the quality of the product he chooses. The production costs for each product quality are displayed in the table below:

Quality	1	2	3	4	5	6	7	8	9	10
Production costs	0	1	2	4	6	8	10	12	15	18

The income of your seller is higher, the higher the price which he is paid. Further, his income is higher, the lower the product quality he supplies to you.

The incomes of all buyers and sellers are determined in the same way. **Each buyer can therefore calculate the income of his seller(s) and each seller can calculate the income of his buyer.** Further, each buyer and seller is informed of the identification number of his trading partner in a given period.

Please note that buyers and sellers can incur losses in each period. These losses have to be paid from your initial endowment of money or from earnings in other periods.

You will be informed about your income and the income of your seller on a separate **"income screen"**. On the screen the following information will be displayed:

- Which seller(s) you traded with
- Which price(s) you paid
- Your desired quality(ies)
- The actual product quality(ies) supplied by your seller(s)
- The income of your seller(s) in this period
- Your income in this period.

Please enter all the information into the documentation sheet supplied to you. After the income screen has been displayed, the respective period is concluded. Thereafter the trading phase of the following period commences. Once you have finished studying the income screen pleas click on the "OK" button.

The sellers also view an income screen which displays the above information. They see the ID of their buyer, the price, desired and actual product quality as well as their own income, your income and—if you have entered two trade agreements—the income of your other seller.

The experiment will not commence until all participants are completely familiar with all procedures. In order to make sure that this is the case we kindly ask you to solve the exercises below.

In addition we will conduct a **trial of the trading phase**, so that you can get accustomed to the computer. This trial phase will not be added to the result of the experiment and therefore not remunerated. Following the trial phase we will begin the experiment which will last for 18 periods.

D.5 Instructions for Chapter 4 (TS Treatment)

Brief overview

Now today's main experiment starts. You already know that your former decisions do not affect your chances at this experiment. Please read the following information carefully. At the end of the information you will find some sample questions. After all participants have answered the questions correctly, this part of today's session will start.

At the beginning of the experiment **groups consisting of 4 participants will be assigned randomly.** Your income in this experiment will depend on your own decisions and the decisions of the other three members of your groups. The experiment will remain anonymous, which means that you will not know at any time who the other group members are.

The experiment consists of **one round** divided into **two stages**. At each stage the participating group members choose a **decision number** independently from each other.

The choice of the decision number is associated with certain **costs.** The sum of your decision number and a number chosen randomly by the computer will determine your **result**. The result of the first stage will determine whether you participate in the second stage: **In the first stage all four group members participate; in the second stage only the two group members with the highest results participate.**

Your income will be calculated in the following way: The group member with the highest result in the second stage will receive the amount of **20.00 Euro** and has to pay the costs of his two decision numbers. The participant with the lower result in the second stage will receive the amount of **12.11 Euro** minus the costs of his two decision numbers. The two participants who did not participate in the second stage will receive the amount of **5.73 Euro** each minus the cost of their decision number in the first stage.

At the end of the second stage, each participant will be informed about his resulting income. After this notification the experiment is over. Your income will be paid to you in cash at the end of the session.

Detailed Information

During this experiment you are in a **group together with three other participants**, so each group consists of four persons. You will not know at any time who the other members of your group are. The experiment consists of exactly **one round** divided into **two stages**. Take your time to think about your decision(s).

Stage 1

In the first stage all participants choose a **decision number** independently from each other. This number can be **between 0 and 125**.

When choosing your decision number you will produce some costs, which you can see in the enclosed cost table. In general, it can be said that the higher the decision number, the higher the cost. The attached **table** contains all **eligible decision numbers and the respective costs**. For example, a decision number "0" causes costs of 0 Euro, a decision number of 50 causes costs of 1.11 Euro, etc.

Your result depends on the chosen **decision number**. Additionally, your result depends on a **random number**. This random number is a number **between -60 and 60** randomly chosen by the computer. Each number between -60 and 60 has the same probability to be selected. The computer will choose the random number **for each participant individually** and independently, which means that the random numbers of the individual participants can be different from each other.

Your decision number (chosen by yourself) and the random number (chosen by the computer) will determine your result for the first stage:

Your result (stage 1) = your decision number + your random number

Your result is higher when your decision number is higher. That means that the **probability to reach the second stage increases with a higher decision number**.

Additionally, your result increases with the random number. For example, if you choose a decision number of 30 and the computer chooses a random number of 15, your result is 45. If you choose a decision number of 93 and the random number is -34, your result is 59.

At the end of the first stage, all participants will be notified about their random number and their result. The result of the first stage determines whether or not you participate in the second stage. Only the participants with the two highest results in your group will participate in the second stage. If two group members have the same result, the participant in the second stage will be drawn by lot. The two group members with the lowest results in the first stage will not participate in the second stage. They receive 5.73 Euro minus the cost of their decision number.

Stage 2

The process in stage 2 resembles the process in stage 1. The two group members who participate in the second stage again get to **choose a decision number**, which is associated with the costs found in the cost table. Additionally, the computer chooses a **new random number** for each of the two participants **individually**. As in stage 1, your result is determined according to the following formula:

Your result (stage 2) = your decision number + your random number

Your result in the first stage thus has **no influence** on your result in the second stage. At the end of the second stage, you will see your decision number and your random number, and you will know if you got the lower or the higher result in this stage. The participant with the lower result in stage 2 receives 12.11 Euro minus the cost of his decision numbers (in both stages). The participant with the higher result receives 20.00 Euro minus the cost of his decision numbers (in both stages).

At the end of the second stage, each participant will be informed about their income determined by their decisions and the random numbers. After this notification the experiment is over.

How is your income calculated?

Your income in this experiment depends on your result and the results of the other group members in stage 1. If you have one of the two highest results in your group, your income additionally depends on your result and the result of the other participant in stage 2.

Three cases can be distinguished:

Case 1: You do not have one of the two highest results in stage 1. In this case you do not participate in stage 2 and your income is calculated in the following way:

Income = 5.73 Euro - cost of your decision number in stage 1

Case 2: You achieved one of the two highest results in stage 1 and the lower result in stage 2. In this case your income is calculated in the following way:

Income = 12.11 Euro - cost of your decision number in stage 1 - cost of your decision number in stage 1

Case 3: You achieved one of the two highest results in stage 1 and the higher result in stage 2. In this case your income is calculated in the following way:

Income = 20.00 Euro - cost of your decision number in stage 1- cost of your decision number in stage 2

In sum, your income is higher, the higher the designated amount of money you are entitled to and the lower your costs. Please note that in cases 2 and 3 your costs incurred in **stages 1 and 2** are relevant for the calculation of your income.

If you have any further questions, please raise your hand and we will come over to answer them.

D.6 Schedule of Effort Costs

Cost Table

Decision Number	Cost in Euro
0	0.000
1	0.000
2	0.002
3	0.00
4	0.01
5	0.01
6	0.02
7	0.02
8	0.03
9	0.04
10	0.04
11	0.05
12	0.06
13	0.08
14	0.09
15	0.10
16	0.11
17	0.13
18	0.14
19	0.16
20	0.18
21	0.20
22	0.22
23	0.24
24	0.26
25	0.28
26	0.30
27	0.32
28	0.35
29	0.37
30	0.40
31	0.43
32	0.46
33	0.48
34	0.51
35	0.54
36	0.58
37	0.61
38	0.64
39	0.68
40	0.71
41	0.75
42	0.78
43	0.82
44	0.86
45	0.90
46	0.94
47	0.98
48	1.02
49	1.07
50	1.11
51	1.16
52	1.20
53	1.25
54	1.30
55	1.34
56	1.39
57	1.44
58	1.50
59	1.55
60	1.60
61	1.65
62	1.71

Decision Number	Cost in Euro
63	1.76
64	1.82
65	1.88
66	1.94
67	2.00
68	2.06
69	2.12
70	2.18
71	2.24
72	2.30
73	2.37
74	2.43
75	2.50
76	2.57
77	2.64
78	2.70
79	2.77
80	2.84
81	2.92
82	2.99
83	3.06
84	3.14
85	3.21
86	3.29
87	3.36
88	3.44
89	3.52
90	3.60
91	3.68
92	3.76
93	3.84
94	3.93
95	4.01
96	4.10
97	4.18
98	4.27
99	4.36
100	4.44
101	4.53
102	4.62
103	4.72
104	4.81
105	4.90
106	4.99
107	5.09
108	5.18
109	5.28
110	5.38
111	5.48
112	5.58
113	5.68
114	5.78
115	5.88
116	5.98
117	6.08
118	6.19
119	6.29
120	6.40
121	6.51
122	6.62
123	6.72
124	6.83
125	6.94

D.7 Elicitation of Risk Attitudes

In the following, we describe the lottery procedure that was used to elicit participants' risk attitudes for the analysis in Chapter 4. The same procedure (with slightly different payoffs) was also used in the experiment of Chapter 1.

Instructions

Before the start of today's main experiment, this first part of the experiment deals with individual decision making. This means that your income in this part only depends on your personal decisions and the outcomes of lottery drawings.

This part of the experiment deals with choices between a **lottery** and a **safe payment**. In the following, you are presented 15 situations. In each situation the lottery is the same, but the safe payment varies. In the lottery you receive 4 Euro with a probability of 50% and 0 Euro with a probability of 50%.

On the following screen you can see the 15 situations. Please decide for each situation if you want to choose the lottery or the safe payment.

At the end of the experiment one of the 15 situations is randomly chosen by the computer. According to your decision in this situation you then either take part in the lottery or you receive the respective safe payment. If you have chosen the lottery in this situation, the outcome of the lottery will also be randomly determined by the computer at the end of the experiment. Your income will be paid to you in cash at the end of the experiment.

Please press CONTINUE to make your decisions.

Lottery Choices

Please decide for every situation whether you choose the lottery or the safe payment!

Situation	Lottery	Safe Payment	Your Decision	
			Lottery	Safe Payment
1	50% Chance 4 Euro and 50% Chance 0 Euro	0.25 Euro		
2	50% Chance 4 Euro and 50% Chance 0 Euro	0.50 Euro		
3	50% Chance 4 Euro and 50% Chance 0 Euro	0.75 Euro		
4	50% Chance 4 Euro and 50% Chance 0 Euro	1.00 Euro		
5	50% Chance 4 Euro and 50% Chance 0 Euro	1.25 Euro		
6	50% Chance 4 Euro and 50% Chance 0 Euro	1.50 Euro		
7	50% Chance 4 Euro and 50% Chance 0 Euro	1.75 Euro		
8	50% Chance 4 Euro and 50% Chance 0 Euro	2.00 Euro		
9	50% Chance 4 Euro and 50% Chance 0 Euro	2.25 Euro		
10	50% Chance 4 Euro and 50% Chance 0 Euro	2.50 Euro		
11	50% Chance 4 Euro and 50% Chance 0 Euro	2.75 Euro		
12	50% Chance 4 Euro and 50% Chance 0 Euro	3.00 Euro		
13	50% Chance 4 Euro and 50% Chance 0 Euro	3.25 Euro		
14	50% Chance 4 Euro and 50% Chance 0 Euro	3.50 Euro		
15	50% Chance 4 Euro and 50% Chance 0 Euro	3.75 Euro		