

---

## The Behavioral Economics of the Labor Market: Central Findings and Their Policy Implications

Ernst Fehr, Lorenz Goette, and Christian Zehnder

Managers claim that workers have so many opportunities to take advantage of employers that it is not wise to depend on coercion and financial incentives alone as motivators. . . . Employers [believe] that other motivators are necessary, which are best thought of as having to do with generosity (Bewley 1995, 252).

Such employment policies are difficult to explain with an economic model in which all agents have strictly selfish preferences. In this paper we argue that these phenomena can be better understood if one acknowledges that a significant share of individuals has a preference for reciprocal fairness that leads them to work harder when they are treated fairly by their employers. We review the evidence on two major psychological forces that possibly drive such fair-minded behavior. First, research suggests that some individuals are willing to sacrifice considerable resources to prevent unfair outcomes: they may be willing to put in extra effort if they feel treated fairly, but they may also withhold effort if they feel treated unfairly (Fehr and Gächter 2000). Second, in making judgments about fairness, individuals compare what they (and others) get to what they think they (and others) are entitled to receive (Kahneman, Knetsch, and Thaler 1986). If any party receives less than this entitlement, fair-minded individuals will try to reduce this gap. A strong feature of these comparisons is loss aversion—that is, the tendency for losses to loom larger than gains.

The presence of these two forces makes specific predictions for the labor-market setting. In particular, high wages should motivate fair-minded workers to exert more effort to reciprocate such fair treatment. A recent literature has produced evidence from laboratory experiments capturing the stylized features of labor markets. The results show that,

indeed, average employee effort is higher when wages are high. Because all the interactions between workers and firms are strictly one-shot, this impulse can only be attributed to a preference for fairness. However, the results also show strong heterogeneity in how individuals respond to receiving high wages. While some employees are fair-minded and reciprocate with high effort, others are unfazed and choose to exert minimal effort, as predicted by selfish preferences.

More recently, field experiments have also been conducted to test this mechanism outside the laboratory. The results are similar to the laboratory experiments: high wages elicit somewhat higher effort, but the effects are generally small. Just like in the lab experiments, it seems that a considerable fraction of individuals is strictly selfish and take the higher wage without reciprocating by expending higher effort.

The behavior of these selfish individuals towards their employers may change, however, in repeated interactions, which more accurately represent most real-life employment relationships. In this more realistic setting, even a strictly selfish individual has an incentive to mimic a fair-minded person and exert high effort when offered a rent because this behavior guarantees high wages in the future. Shirking, in contrast, reveals that the individual is egoistic. Firms will be unwilling to pay wage rents to selfish workers once they have proven that they don't exert much effort. Thus, finite repetition can greatly magnify the effect of fairness on labor market outcomes. This mechanism can be tested explicitly in experiments and proves to be very powerful. We also discuss evidence from several field studies in which employers violated fairness norms in ongoing relationships with their employees. The results in these studies show that employees respond strongly to how they are treated.

What predictions does a fairness model make for wage dynamics? Given the evidence on how fairness is judged, the behavioral model of wage dynamics predicts that entry-level wages will be cyclical, while the wages of incumbent workers should be largely unresponsive to changes in labor-market conditions. Moreover, loss aversion implies that wage cuts should be particularly rare. We examine the experimental and field evidence on these issues and find the model's predictions largely confirmed in the data. The wages of newly hired workers are highly cyclical, while the wages of incumbent workers hardly respond to the business cycle. Further, there is

strong evidence that wages are downwardly rigid. In many cases, firms prefer to freeze workers' wages rather than cut their wages.

We then discuss several policy implications that emerge from the model, as well as the support for these prescriptions in the data. We consider two important aspects of economic policy—namely, monetary policy and minimum wages. In both cases, the psychological forces underlying the behavioral model of the labor market lead to important new conclusions. They also highlight the importance of forming a detailed understanding of the underlying psychological motivations of any outcome, as the policy implications may differ strongly depending on the model that generates the outcomes. Perhaps our strongest conclusion from the research that we survey here is to caution against “rationalizing”; that is, tweaking the economic model along “standard” dimensions to make it fit the data. We argue that such an approach can be highly misleading and lead to bad policies. Rather, we suggest that microfounding a model and testing each ingredient is a far more fruitful approach for the purpose of informing policy.

The rest of this study is structured as follows: section 1 summarizes the psychological forces central to understanding several labor market outcomes. Section 2 explains how these forces change the tradeoffs that firms face and summarizes the evidence on these tradeoffs. Section 3 illustrates how these findings can be relevant for policy. Section 4 concludes by summarizing how insights from behavioral economics can foster a better understanding of how the labor market functions, and how policies can be designed to build on these insights.

## 1. Psychological Forces

Standard economic analysis assumes that individuals' preferences are defined over their own consumption only, and that consumption enters the utility function in levels. However, evidence from two decades of research in experimental economics has shown this assumption to be wrong. In this section we discuss two of the most intensively discussed topics of this literature. The first concerns the assumption of selfishness, meaning that just one's own consumption enters the utility function. The evidence overwhelmingly shows that individuals also care about other

people's consumption. The second concerns the way in which consumption enters the utility function. While the standard economic model assumes that only the level of consumption matters, the evidence shows that the level of consumption is valued relative to a reference level. As we will discuss in section 2, incorporating these two features of preferences leads to a new understanding of important labor market phenomena.

### *Testing Self-Interest: Ultimatum Games*

Striking evidence against the self-interest hypothesis comes from a large number of experiments which study the ultimatum game. In this game, introduced by Güth, Schmittberger, and Schwarze (1982), two players have to agree on how to divide a fixed sum of money. The game is structured as follows: At the beginning of the game, one player (the proposer) gets a money endowment. The proposer suggests how the endowment should be divided between him and the second player, the responder. To this end, the proposer makes a take-it-or-leave-it offer to the responder. Then the responder may only decide whether to accept or reject the proposer's offer. If the proposed split is accepted, the money is divided according to the proposer's offer. If the offer is rejected, both players receive nothing.

The self-interested model in standard economics makes a distinct prediction for the ultimatum game's outcome: the proposer offers the smallest positive amount of money, say \$1, and the responder accepts this offer. However, in strong contrast to this self-interest prediction, many experimental studies show that subjects are willing to reject unfair offers, even if this rejection is associated with substantial costs. A robust result across a large number of studies is that offers of less than 20 percent of the available surplus are rejected with a probability of 0.4 to 0.6. The probability of rejection decreases the larger the size of the offer. Once offers reach 40 to 50 percent of the available surplus, they are only very rarely rejected (see, for example, Camerer and Thaler 1995; Fehr and Schmidt 2003; Güth, Schmittberger, and Schwarze 1982; and Roth 1995). Rejections of low offers in the ultimatum games show that many subjects do not care solely about their own payoff. A plausible interpretation for the rejection of low offers is that subjects perceive them as unfair. Receiving the offer of only a small share when the proposer could have chosen a

more equitable split seemingly lowers the responder's utility. Yet if the utility-reducing impact of unfairness is strong enough, the responder may prefer to forgo his share in order to avoid accepting the unfair outcome.

Despite the large number of replications, the experimental findings in ultimatum games have often been contested. A very frequent objection is that the usual monetary amounts used in experimental economics are so small that people do not really care about the decisions they make in the experiments. However, several studies suggest that even very large increases in stakes (up to several months' wages) have surprisingly modest effects on behavior in ultimatum games (see, for example, Cameron 1999; List and Cherry 2000; and Slonim and Roth 1998). The data indicate that when the stakes are high, responders are slightly more reluctant to reject what they deem as unfair offers. But the fear of costly rejection motivates many proposers to make offers close to 50 percent anyway. Hence, increasing the monetary stakes does not destroy the strong impact of fairness considerations on bargaining outcomes.

The results of ultimatum games are also remarkably robust across different cultures. Studies that compare ultimatum behavior across industrialized countries find only relatively small differences in average offers and rejection rates. Although there seems to be some cross-country heterogeneity in the perception of what constitutes a reasonable offer in such an experiment, there is not a single industrialized country where the outcomes are even close to the self-interest prediction (see, for example, Buchan, Croson, and Johnson 1998; Roth, Prasnikar, Okuno-Fujiwara, and Zamir 1991). However, the impact of culture is much more pronounced when nonindustrialized societies also are investigated. A group of anthropologists and economists have studied ultimatum game behavior in a number of small-scale societies in Africa, Asia, and South America (Henrich et al. 2001, 2004). While many of these cultures exhibit behavior similar to the Western students who are the main subjects of these laboratory experiments, some societies follow a completely different behavioral pattern. For example, the Machiguenga in Peru make much lower offers and have much lower rejection rates than are usually observed in ultimatum games. This behavior is probably due to the fact that the Machiguenga mainly interact within their own families,

while transactions with others (even within a village) are very rare. In general, the comparison of the small-scale societies reveals that the stronger the cooperative activity (for example, collective hunting) and the higher the market integration (common language, trade, and developed labor markets), the more likely it is that a society's sharing norms will be close to equal splits. Hence, in contrast to the standard economic model, in which self-interest and functioning markets go nicely together, these studies reveal that people in societies in which markets play a bigger role behave less rather than more self-interestedly (see Camerer 2003 for a more detailed discussion).

While there is no longer much disagreement regarding the fact that people's behavior in the ultimatum game systematically deviates from the prediction of the rational self-interest model, researchers still somewhat disagree on how this evidence should be interpreted. It is probably fair to say that the majority of experimental and behavioral economists is convinced that the existence of fairness preferences is the most plausible explanation for the ultimatum game evidence.

However, some researchers prefer to interpret the rejections of low offers as a consequence of bounded rationality. Binmore, Gale, and Samuelson (1995) and Roth and Erev (1995) suggest a learning explanation grounded in self-interest. In both arguments the basic idea is that the speed of learning depends on the cost of the error. Since rejecting low offers is not very costly, responders only learn slowly that this choice is not optimal. In contrast, because rejections are very costly to the proposers, they quickly realize that low offers are not profitable. Therefore, this asymmetry in learning makes convergence to the equilibrium very slow. Indeed, simulations in these papers show that thousands of iterations may be necessary to bring behavior close to the standard prediction. While we certainly agree that learning models are useful tools to understand behavioral changes in complicated dynamic setups, we doubt that these models are adequate to describe the behavior in simple environments like the ultimatum game. The responders' decision is so straightforward that it seems hard to believe that they systematically fail to make the optimal choice dictated by their preferences.

Other economists put forward the idea that the behavior in ultimatum games may be a consequence of the erroneous application of repeated

game behavior to one-shot laboratory environments (for instance, see Binmore 1998 and Baker in this volume). The argument is that in everyday life the vast majority of decisions are made in the context of repeated interactions. It is well-known that the rejection of low offers can be sustained in equilibrium in an infinitely repeated game of agents with purely selfish preferences. Since people are used to repeated interactions, the appropriate behavior therefore may evolve into a social norm, which people then also apply in the somewhat unnatural one-shot situations in the laboratory. One problem with this explanation is that the folk theorem tells us that, in general, repeated games are characterized by a plethora of equilibria (for details, see Rubinstein 1979). In some cases, the different equilibria can be Pareto-ranked such that efficiency arguments yield equilibria that may be considered focal. In other games, however, such refinements do not work: in the ultimatum game, for example, every outcome in which the responder accepts the proposer's offer can be sustained as an efficient equilibrium of an infinitely repeated game. Thus, while an equal split is a possible outcome of the repeated game, so is any other division of the pie. Accordingly, this alternative approach is not able to explain why people tend to share the pie equally in the ultimatum game.

On the other hand, we agree that everyday experience affects behavior in the laboratory. This is evident from the results of the cross-cultural experiments in small-scale societies: groups who strongly depend on cooperative behavior are more likely to share the pie equally in experiments. However, the fact that external social norms affect people's behavior does not imply that they are unable to distinguish one-shot games from repeated setups. In contrast, the available evidence shows that the possibility of acquiring a reputation or repeatedly meeting the same partner strongly increases people's willingness to cooperate and to punish uncooperative behavior. For example, Fehr and Fischbacher (2003) show that responders in an ultimatum game have a significantly higher acceptance threshold if future proposers are informed about their current rejection behavior than if future proposers do not get information about the responders' past behavior (for more laboratory evidence on the impact of reputation formation on behavior, see Andreoni and Miller 1993; Brown, Falk, and Fehr 2004; Brown and Zehnder 2007; Engelmann and Fischbacher 2009; Fehr and Zehnder 2009; Gächter and Falk

2002; and Seinen and Schram 2006). In our view, the fact that changes in strategic incentives strongly affect people's behavior makes it unlikely that the deviations from selfish behavior in one-shot games are merely the consequence of mistakenly applied rules of thumb. It seems much more plausible that the prevailing social norms in a society directly affect their preferences.

Furthermore, recent evidence from neuroeconomics provides additional support for the assumption that people have a preference for fairness. For instance, the human brain's reward center is activated if people can punish others that defected against them in a game (De Quervain et al. 2004). Singer et al. (2006) conduct a study in which a confederate cooperates or defects against the subjects in a first experiment. In the second experiment, the confederate and the subjects receive moderate electric shocks. If the confederate defected in the first experiment, the reward center of the subject activates when the confederate receives an electric shock. These results show that there is a direct link from experiencing unfair behavior against oneself to gaining utility by retaliating.

The combination of these arguments makes us think that the assumption that a considerable share of people prefer fair outcomes is the most plausible explanation for the observed behavior in ultimatum games and in related experiments.

### *The Reference Frame for Judging Fairness*

In laboratory experiments like the ultimatum game, the equal split of a surplus seems to serve as a natural reference transaction by which the fairness of an outcome is evaluated. However, it is rather unrealistic to assume that this finding can be extended to other settings. Thus, in general, the determination of the fairness standard or the reference transaction is likely to be dependent on the specific environment and context. Kahneman, Knetsch, and Thaler (1986) use questionnaires to investigate how standards of fairness are determined. Their study presented participants with a number of different business scenarios. In each scenario, a firm either lowered wages or increased prices in response to an external shock. For each scenario, the participants had to indicate whether they perceived the wage (price) change to be acceptable or unfair. The responses of participants confirm the hypothesis that the perceived fairness of an action is also dependent on the specific context. In one question,

for example, they asked the participants whether they found it acceptable for a firm to lower a current employee's wage after an increase in unemployment had enabled other firms to hire similar workers at lower wages. A large majority of people perceived this action as unfair. Here, it seems that the fairness standard is determined by the past interactions between the worker and the firm. If in the same situation, however, the current worker left and the firm hired a new worker at a lower wage, the majority of people found this acceptable. In this case, the new worker and the firm do not have a common history, and, accordingly, the interactions of other workers with other firms serve as the reference point for evaluating fairness. This example shows that past interactions provide a context that has a strong impact on the reference transaction used to judge the fairness of a particular action: identical outcomes may trigger very different fairness judgments if they take place in different contexts.

Falk, Fehr, and Fischbacher (2003) provide experimental evidence on how the economic environment can affect fairness judgments by conducting a series of ultimatum games in which they restrict the proposer's strategy set. In one treatment, the proposer has an endowment of ten and can offer either two or five to the responder. In the second treatment, the proposer can offer either zero or two from the endowment of ten. Thus, in both treatments, there is an offer that gives the responder two, while the proposer gets eight. Their results show that the offer of two is more frequently rejected when the proposer could have offered five than when the proposer could have offered zero.

This finding suggests that the set of available actions determines the fairness standard. When the alternative is an offer that would have given five to both subjects, the offer of two is often rejected because it is perceived as unfairly low. However, when the only alternative is the offer of zero, then the offer of two is considered kind, and, accordingly, the rejection probability is low. The behavioral evidence provided in this study confirms that not only the chosen action and its consequences matter—because in different situations the same action may reveal completely different underlying intentions. Thus, the fairness of a particular action is often not only determined by the resulting payoffs but also by the set of available, yet unchosen, alternatives.

In the context of labor markets, an enduring and important question is whether individuals only take into account the real buying power of

their wages, or whether they also care about nominal wages. Dating back at least to Keynes (1936), some economists have assumed that individuals are not fully aware of changes in the price level, and thus also care about the nominal wage. In the context of fairness judgments, Shafir, Diamond, and Tversky (1997) have demonstrated that individuals understand how inflation changes their purchasing power, and correctly take this factor into account when judging economic welfare in hypothetical scenarios. However, Shafir, Diamond, and Tversky (1997) consistently show that individuals also believe that higher *nominal* wage growth leads to higher happiness for other individuals, suggesting a strong belief that other individuals are unable to back out the inflation rate correctly. Consistent with the first set of results in this study, other studies of consumers' inflation expectations have also shown these to be surprisingly accurate predictions (see Mankiw, Reis, and Wolfers 2004). This evidence seems to suggest that individuals may well be able to anticipate price changes, but that they often believe that others cannot.<sup>1</sup> Goette and Huffman (2007b) show that for fairness judgments, individuals only care about real wage changes as long as the nominal wage is increasing. However, if a nominal wage cut occurs, this is viewed as distinctly more unfair. Goette and Huffman argue that the salience of a nominal wage cut triggers strong negative emotions that sway the fairness judgment: because a wage cut is perceived as unambiguously bad, this event leads to a much stronger reaction than when the nominal wage is rising but more slowly than the price level. Therefore, this evidence suggests that for fairness judgments, individuals do take into account the price level and effectively care about real wage changes. However, if the nominal wage were to be cut, this leads to a stronger reaction.

Wage cuts, in general, seem to be perceived as much more unfair than the same reduction in pay if these cuts are perceived as eliminating a gain. For example, Kahneman, Knetsch, and Thaler (1986) gave respondents a scenario where a firm used to pay a 10 percent bonus every year, but then abolished this practice. The vast majority of respondents considered this action to be fair, even though it effectively cut the workers' incomes by 10 percent. Other respondents were given a scenario where the workers base wage was cut by 10 percent. In this case, the majority of the respondents stated that the action was unfair.

### *Reference-Dependent Utility*

The previous examples showed that in an economic environment fairness judgments depend upon reference transactions. This property is not limited to the realm of fairness judgments. In a seminal paper, Kahneman and Tversky (1979) argue that utility depends not only on the level of consumption, but on reference points to which consumption is compared. Kahneman and Tversky propose two central features of how individuals value outcomes relative to reference points: (i) *loss aversion*: falling short of the reference point by one unit hurts more than exceeding it by one unit is pleasurable; (ii) *diminishing sensitivity*: Kahneman and Tversky argue that the marginal benefit decreases with the distance from the reference point.

Loss aversion provides a rationale for why individuals exhibit risk aversion over small stakes and buy, for instance, home telephone-wire insurance or extended warranties. As pointed out in Rabin (2000) and Rabin and Thaler (2001), expected utility theory cannot accommodate risk aversion over such small stakes without making predictions that seem outright crazy. For example, most individuals would reject a coin flip in which they could win \$110 or lose \$100.<sup>2</sup> If an expected-utility maximizer rejects this coin flip for any wealth level between her current status quo and, say, \$100,000 more, then she would be unwilling to accept a 50:50 gamble in which she can win \$100,000 or lose \$220, irrespective of the specific shape of the individual's utility function.<sup>3</sup> Such behavior, which is implied if expected utility is to account for the rejection of small-stake gambles, strikes us as implausible, and makes it clear that something else is needed to explain why individuals reject such gambles. If individuals compare outcomes relative to a reference point such as the status quo, it is easy to see how loss aversion can lead an individual to reject the small-stakes gamble while avoiding the implausible implications for larger stakes.

Loss aversion can also help explain a behavioral pattern in a different area of research. In a classic study by Kahneman, Knetsch, and Thaler (1990), individuals value a good more highly when they have to give it up than when they can acquire it. This is consistent with the interpretation that individuals perceive having to give up an object they expected



to keep as a loss, while they perceive acquiring the same object as a gain, and spending the money to obtain it as a loss. Therefore, loss aversion predicts that selling reservation prices should be higher than buying reservation prices. A recent study by Gächter, Johnson, and Herrmann (2007) measures subjects' loss aversion over risky gambles and the buying and selling prices for model cars. Indeed, they find that more risk-averse individuals tend to display a disparity between the prices at which they are willing to buy and to sell.

The evidence with respect to diminishing sensitivity is more mixed. Diminishing sensitivity predicts increasing marginal utility towards the reference point; that is, it predicts a concave valuation function over gains, but a convex valuation over losses. This implies that individuals should be more willing to gamble when deciding between a sure loss or an unfair lottery offering the chance of not incurring any loss. Kahneman and Tversky (1979) found strong evidence of this pattern, but many of the scenarios were hypothetical. Several studies have now shown that the incidence of risk-seeking over losses is smaller when the decisions involve real payoffs (Holt and Laury 2002).

Recent evidence has renewed researchers' interest in diminishing sensitivity. Post et al. (2008) examined contestants' behavior on the game show *Deal or No Deal*. Some lucky contestants are virtually certain to win large amounts of money, while it quickly becomes clear to others that it is very unlikely they will win much. There is a strong shift in the risk preferences of the contestants who have bad luck: their strategies become risk-seeking, while the strategies of the lucky contestants become highly risk-averse. These preferences are consistent with the interpretation that unlucky contestants face losses relative to what they could have expected from the game, and their diminishing sensitivity to these likely potential losses makes them risk-seeking in an effort to reverse their luck for the better.

The principles of valuation laid out in prospect theory seem to apply to a wide range of phenomena. For example, Heath, Larrick, and Wu (1999) report that (arbitrarily set) goals (for example, "do 50 push-ups") assume the properties of a reference point. Falling short of a goal by one unit hurts more than the pleasure of exceeding it by one unit. An exten-

sive literature also documents the goal-gradient effect. The goal-gradient effect means that an individual will exert more effort to reach a goal the closer he is to the goal. The effect can easily be explained by diminishing sensitivity and loss aversion: if goals inherit the properties of reference points, then the closer an individual is to the goal the steeper the value function, and hence the higher the marginal utility gained from making progress towards the goal. Once the goal is achieved, making additional progress only feels like a small gain and is consequently worth less.

Reference-dependent preferences can have important implications in labor markets, which we detail in the next section. They also have important implications in other areas, as recent research has shown. Odean (1999) shows that investors are much more likely to hold on to stocks with paper losses in their portfolio, despite the strong tax incentives to sell them. Similarly, Genosove and Mayer (2001) find that a homeowner facing a loss (relative to the home's purchase price) asks a higher selling price, holding all other characteristics of the house constant. They also find evidence of diminishing sensitivity: at the margin, a small loss increases the asking price proportionally more than a large loss. Gächter, Johnson, and Herrmann (2007) demonstrate the effect of reference-dependent preferences when evaluating a product's attributes: prospective car buyers require a much larger reduction in a car's purchase price if a feature is removed than they are willing to pay to have the feature added. This disparity again correlates with behavior in lotteries with small-stake risks, lending further credence to the interpretation that loss aversion is driving these disparities.<sup>4</sup>

## 2. A Behavioral View of the Labor Market

Firms invest significant resources to treat their employees well (see, for example, Bewley 1999): they pay high wages to elicit high effort, refrain from cutting wages because they fear the negative consequences for employee motivation, and provide various other benefits in the hopes of keeping their employees satisfied. However, in the standard model in economics, firms need not care about fairness. In fact, as we argue below, even in the context of repeated interactions, treating employees well

makes no business sense. The reason is that, ultimately, any interaction an employee has with a firm is finite. Paired with *strictly* selfish preferences, this makes repeated-game incentives typically ineffective.

However, the psychological forces we outlined in section 1 can alter the calculus of firms dramatically. Even in one-shot interactions, fair treatment by firms may motivate workers to exert more effort. Further, as we explain, these effects are amplified when combined with *finite* repetition, as it is plausible in many circumstances.

In this section, we outline the predictions of an economic model of the labor market that incorporates the psychological forces reviewed in section 1. We argue that key aspects of labor markets can be better understood if one takes these forces into account. For example, phenomena such as the very rare occurrence of nominal wage cuts, or the observation that employees do not respond to incentives as predicted by the standard theory, can be explained if concerns for fairness and loss aversion are taken into account.

### *Characteristics of Employment Relationships*

A central observation is that the typical employment relationship is characterized by incomplete contractual agreements and repeated interactions. Contractual incompleteness means that the legal agreement between the firm and its workers does not determine many relevant details of the workers' jobs. The reason for the incompleteness of employment agreements is that most occupations consist of multidimensional and complex tasks that can neither be completely foreseen nor perfectly described. Incomplete contracts imply that outsiders can hardly determine whether the trading parties have met their obligations or not. As a consequence, important aspects of the collaboration of firms and workers cannot be enforced by third parties. This lack of enforceability obviously creates a fundamental problem for the firm: if its employees' duties and obligations are only vaguely specified, how can a firm motivate its workers to provide more than minimal effort?

Sometimes the answer to this problem is straightforward. In simple jobs, it may well be that a worker's output is relatively easy to measure. In these cases the firm can use explicit incentive contracts to motivate its workers. If a worker's earnings depend on his output, it may be in his

interest to exert nonminimal effort, even though nobody can force him to do so.<sup>5</sup> In many jobs, however, output is complex and while some dimensions are objectively measurable, others are not. In all these cases, the provision of explicit incentives may lead to distorted outcomes because workers will allocate all of their effort toward those activities that are rewarded by the firm's incentive scheme. Thus, if workers are expected to devote time and effort to activities which contribute to the nonmeasurable dimensions of output, incentive pay cannot be effectively used. In the literature, this so-called "multi-tasking problem" has been emphasized as one of the main reasons why employment contracts often stipulate a fixed wage payment (see, for example, Baker 1992; Holmstrom and Milgrom 1991).

However, several researchers have argued that even if firms are constrained to paying fixed wages, they can make use of the long-term nature of employment relationships to force high levels of effort from workers. If firms and workers have the possibility of interacting repeatedly, firms can condition the future terms of their employment contracts on workers' current performance. This incentive can motivate workers to provide high effort because the firm will only continue to pay a high wage in the future if they exert the desired effort today (see Bull 1987; Hart and Holmstrom 1987; MacLeod and Malcomson 1989, 1993, 1998; Shapiro and Stiglitz 1984). While intuitively appealing, this argument has a serious flaw: as long as all market participants have completely selfish preferences, it requires infinite repetition. If either firms or workers have a finite time horizon, backward induction implies that performance in employment relationships collapses to the minimal required effort. For firms, the assumption of an infinite time horizon can be justified. Even though firm owners have finite durations, they may be interested in maximizing their firm's long-term value because this value determines the revenue they will receive when they sell their assets to a successor after retirement. The assumption of infinitely lived workers, in contrast, is problematic. Employees retire from the work force with certainty. Thus, from the perspective of workers, the relevant duration of an employment relationship may be long, but always finite. But this sets in motion a process of unraveling: as soon as a worker is close to retirement, the firm can no longer threaten to lower wages in the future (because there is no future). As a



consequence, there is no way for the firm to keep a selfish worker from shirking shortly before retirement. However, since the firm anticipates the worker's behavior, the firm has no incentive to pay a high wage in the period before retirement. This, in turn, destroys the employee's incentives to work in the second-to-last period. This argument can be iterated back to the beginning of the employment relationship. Consequently, finite repetition with selfish workers cannot motivate above-minimum effort.

The prediction that repeated interactions do not change labor market outcomes seems counterintuitive in light of everyday experience. Indeed, we argue that repeated-game incentives in the labor market are effective *because* some individuals care about fairness. In this case, paying a high wage can elicit high effort even in a one-shot interaction. It is the preference for fairness that leads a fair-minded worker to choose a higher effort to raise the firm's payoff because the firm just raised his payoff (Akerlof 1982; Akerlof and Yellen 1988, 1990; see Benjamin 2006 for a formal model). Yet while many individuals may have such strong preferences for fairness, the evidence reviewed in section 1 also suggests that a considerable fraction of individuals is more or less selfish. As a consequence, the one-shot effects of fairness are not always strong enough to make it profitable for firms to pay noncompetitive wage premia.

However, only few people with fairness concerns are needed in order for repeated interaction to have a powerful impact. The results of the seminal paper by Kreps et al. (1982) imply that the presence of a small fraction of nonselfish workers can give perfectly selfish workers an incentive to work hard, even though the worker and firm interact only a finite number of times (see Brown, Falk, and Fehr 2004 for an example). The role of repetition in this case, though, is different from the one in games with an infinite horizon. The intuition is that selfish agents now have an incentive to exert effort in order to maintain a reputation and make the firm believe that they are (at least potentially) fair-minded. Such a reputation is valuable for selfish workers because finite repetition implies that the firm only pays noncompetitive rents to workers who have not yet been identified as selfish. The firm anticipates that a selfish worker will always shirk in the final period of the interaction, which unravels all incentives to pay a rent in any period to a worker known to be selfish. Fair-minded workers, in contrast, exert effort whenever they are paid a fair wage

involving a rent. Thus, if the belief about a worker's fair-mindedness is high enough, the firm will be willing to pay a high wage even in the last period: even though the worker may turn out to be selfish sometimes, *ex ante* it may still be profitable to pay a rent in the last period because such a wage will still cause the worker to exert effort if he is fair-minded. Note that it is the payment of a rent that disciplines the selfish worker, because if the worker shirks he will be identified as a selfish type which, in turn, implies that he will not be paid a rent in the future.<sup>6</sup>

### *Gift Exchange in One-Shot Interactions*

For the above argument of long-run reputation building to work, a necessary condition is that there is some gift exchange in one-shot relationships; that is, there are at least some individuals who respond to higher wages—or better treatment more generally—by exerting more effort. Therefore we examine the evidence from one-shot employment relationships in this subsection.

The cleanest test of gift exchange comes from laboratory experiments that contain essential strategic features of an employment relationship. As we will show in the first part of this section, the laboratory studies provide clear evidence that a nonnegligible number of individuals are fair-minded and respond to receiving higher wages by exerting more effort. However, the evidence also shows that fairness concerns alone may not be sufficient to make it profitable to pay noncompetitive wage rents. In fact, if fairness concerns are the only force that drives worker effort above nonminimal levels, there are (in general) large unexploited efficiency gains.

In interpreting the lab studies, there is a natural concern that the effects of gift exchange found in the lab may not carry over to the real world (Levitt and List 2007). In particular, because the environment in the lab experiments is intentionally stylized, it is not obvious whether the effects of fairness may also be measurable in labor markets outside of the laboratory. In the second part of this section we discuss a number of field experiments which explicitly manipulate the wages paid to workers in real-life work environments. In general, these studies suggest that the effects of wage changes on workers' productivity may also be present in the real world, even in setups characterized as one-shot interactions.

*Evidence from the Lab*

Laboratory evidence for the positive impact of wage rents on the amount of effort a worker exerts come from studies of the so-called gift exchange game, which has been introduced by Fehr, Kirchsteiger, and Riedl (1993). The gift exchange game is a two-player game that captures the basic features of a principal-agent relationship with highly incomplete contracts. The structure of the game is as follows: First, the employer offers a wage which the worker can accept or reject. In the case of acceptance, the worker must choose an effort level which cannot be enforced by the employer. If the worker rejects the offer, he receives an unemployment benefit. The firm's profit is equal to the returns generated by the worker's effort minus the cost of this effort. The worker's payoff, in turn, is calculated as the wage minus the effort's cost. The parameters are chosen in such a way that the efficient outcome is achieved when the worker chooses the maximal effort level, but because the provision of higher effort is associated with additional costs, the self-interest model predicts that the worker always provides the smallest possible effort.

However, if fairness considerations matter and workers are willing to reciprocate generous wage offers by providing higher effort levels, it may be profitable for employers to induce a gift exchange by offering wages that exceed the worker's outside option. It is interesting to consider the extreme case of workers with a very strong preference for fairness.<sup>7</sup> In the model used in Benjamin (2005), workers try to equalize the surpluses (the material payoffs minus the reference payoffs they feel are appropriate) they and the firm receive. Thus, if the firm increases their wage by \$1, the workers need to close a \$2 gap in surplus. It is interesting to ask what wage a profit-maximizing firm will set in this context. It turns out that the firm will set a wage that induces the workers to choose the efficient effort level. The key idea behind this insight is understanding how the workers close the \$2 surplus gap after a \$1 increase in their wage. As long as the marginal cost of their effort is below the marginal product of effort supplied to the firm, the firm's revenues will rise by more than \$1, while the workers' cost will rise by less than \$1 in closing the gap. Hence, a \$1 wage increase raises profits. This increase is profitable for the firm up to the point where the marginal cost of effort equals the marginal revenue product of effort, which coincides with efficiency

(see Benjamin 2006 for an extensive discussion and a formal proof). In the case of the gift exchange game discussed here, the prediction is that firms set their wage offer such that it induces the maximal effort level.

A typical example for a one-shot gift exchange experiment is the baseline treatment in Brown, Falk, and Fehr (2004). They implement a simple labor market that has an excess supply of workers. In every period each employer can at most hire one worker and each worker can at most have one job. Although the market runs for several periods, reputation formation cannot play a role because employers cannot distinguish workers from each other when they make wage offers.

The main results of this treatment can be summarized as follows: in line with the wage-effort hypothesis, a worker's effort indeed depends on the firm's wage offer. On average, if firms offer higher wages, workers will provide more effort. As a consequence wages and efforts are higher than predicted by the self-interest model. On average, efforts settle at a level of about 3 (on a scale of 1 to 10). However, while the effort level is significantly higher than the minimal effort of 1 (which the self-interest hypothesis would predict), it is also far from the efficient effort level of 10 (which the complete fairness model would predict). The reason why gift exchange does not have a more strongly positive effect on market performance is due to the huge inter-individual differences across subjects. While there is a considerable fraction of workers whose effort provision exhibits a strongly reciprocal pattern, there is also a substantial fraction of workers who often make purely selfish choices. The relationship between wages and effort is steep enough to render nonminimal wage offers profitable, but the presence of selfish agents restrains many principals from making wage offers which would be high enough to induce efficient effort levels from fair-minded workers. Thus, on the one hand, the evidence in this study confirms the empirical relevance of the gift-exchange hypothesis because, on average, worker effort depends positively on wages. On the other hand, the study also reveals that the impact of one-shot gift exchange interactions on aggregate market efficiency may be small relative to the first-best solution.

The finding that gift exchange has a positive but small impact on market performance is very robust and has also been found in a number of

**Table 1**  
Gift Exchange in One-Shot Interactions: The Laboratory Evidence

Study	Average Effort	Effort Range	Number of Periods	Students as Subjects	Excess Supply of Workers
Fehr, K, and R (1993)	0.40	0.1 -1.0	12	Yes	Yes
Fehr and Gächter (1998)	0.44	0.1 -1.0	12	Yes	Yes
Fehr et al. (1998)	0.37	0.1 -1.0	10	No	No
	0.40	0.1 -1.0	10	No	Yes
Gächter and Falk (2002)	0.41	0.1 -1.0	10	Yes	No
Hannan et al. (2002)	??	0.1 -1.0	12	Yes	Yes
Charness (2004)	0.31	0.1 -1.0	10	Yes	No
Charness, F, and K (2004)	0.32	0.1 -1.0	10	Yes	No
	0.23	0.1 -1.0	10	Yes	No
Brown, F, and F (2004)	3.30	1 -10	15	Yes	Yes
List (2006)	3.50	1 -10	5	No	No
	2.30	1 -5	5	No	No
	2.50	1 -5	1	No	No

*Source:* Authors' calculations.

other laboratory studies using students as participants (see table 1 for an overview). Charness (2004); Charness, Frechette, and Kagel (2004); Gächter and Falk (2002); Fehr and Gächter (1998); Fehr, Kirchsteiger, and Riedl (1993); and Hannan, Kagel, and Moser (2002) all report evidence from one-shot gift exchange experiments and find results which are very similar to the ones reported above: wages and effort levels are always positively correlated, but the realized effort level is far from efficient. Fehr et al. (1998) and List (2006) replicate the gift exchange findings in laboratory experiments with nonstudent subject pools (soldiers and sports card enthusiasts, respectively) and find that the realized effort levels are in line with the evidence presented in the previous literature.

Despite the fact that fairness concerns alone leave a large part of the available welfare gains unexploited, it should be emphasized that gift exchange is a remarkably robust finding even in competitive environments. In many of the studies listed above there is an excess supply of workers, such that workers compete for jobs (see table 1). Nevertheless,

firms are willing to pay noncompetitive rents to workers in order to elicit nonminimal effort levels. Fehr and Falk (1999) test the robustness of this result in yet another way. They induce a labor market with excess labor supply in which both firms and workers can make wage proposals in a double auction. In many cases, this institution is known to quickly converge to the competitive equilibrium predicted by the standard model. They find that firms do not hire workers who try to underbid the going wage because they fear that they will attract shirkers. The unemployed workers try very hard to gain employment by offering to work for very low wages, but the firms prefer to pay higher wages, as, on average, this guarantees them higher effort. Thus, gift exchange prevails even in highly competitive environments.

#### *Evidence from the Field*

Many years ago social psychologists were testing whether generating feelings of being overpaid or underpaid affected work effort (for a review, see Goodman and Friedman 1971). For example, Pritchard, Dunnette, and Gorgenson (1972) attempt to manipulate these perceptions directly by giving (mostly deceptive) information about whether the wage they were paying was considered high or low. The results are generally supportive of the view that if individuals feel treated generously (meaning if they feel overpaid), they exert more effort. However, few economists would accept this evidence, as it relies heavily on deceiving the subjects (see Roth 1995 for a discussion of the use of deception in experiments). Several recent studies explicitly manipulate the wages paid to real-life workers in order to assess whether a higher wage translates into higher effort, much like it does in the experiments. The idea is to test the proposition that higher wages are perceived as more fair and, consequently elicit higher effort. In a pioneering study, Gneezy and List (2006) hire workers to enter books into a library information system. The workers are made aware that this is a one-time employment situation (as once the books are entered, there is no further work for them). The workers are paid hourly, either \$12 (in the baseline condition) or \$20 (in the gift exchange condition), with no particular reason given for the pay rate. Based on the theory reviewed above and the lab experiments, the prediction is that effort will increase when the subjects are paid \$20. Overall, output is approximately 10 percent higher, but this difference is not significant. Thus, while the point

estimate is sizable, it is not significant because of the variation in individual output levels and because of the small sample sizes used in this study. More alarmingly, the authors find a decline in the treatment effect after only a very few hours. Their interpretation is that gift exchange in real life erodes after a very brief period, and effort falls back to its original level. However, it is difficult to back up this interpretation with strong evidence because their sample size is so small.

In a rich study, Al-Ubaydli et al. (2008) correct this problem: their samples are larger (30 subjects per condition, compared to approximately 10 in Gneezy and List 2006). The cleanest comparison in the study is the comparison of the following two conditions. All workers were recruited through a temporary employment agency and given only a vague hourly wage band (\$8 to \$16) that they would be paid for the work that they did stuffing envelopes.<sup>8</sup> Subsequently, some workers were then paid \$8, while others were paid \$16. Again, nothing was said about why they were paid this wage rate, but presumably the individuals who were paid \$8 per hour felt treated less fairly than individuals paid \$16 per hour. There is a large and highly significant output difference between the two treatments ( $p < 0.01$ ): when paid \$8 per hour instead of \$16 per hour, the number of finished envelopes decreased by about 22 percent. Further, this difference showed no sign of declining—if anything, it increased with time. Clearly, the fixed wage rate affected performance in this treatment. However, it is unclear whether the effect was mainly caused by workers feeling that they were treated unfairly after their wage band was announced and they received \$8 an hour instead of \$16, or whether the highly paid workers put in extra effort.

Thus, while it is quite clear that the fixed wage rate affects worker effort, the main treatment in Al-Ubaydli et al. (2008) makes it difficult to disentangle how much of the effect was caused because the workers in the \$8 treatment felt they got less than they were entitled to receive, or because the workers in the \$16 treatment felt they received more compensation. Another study by Kube, Maréchal, and Puppe (2008) provides a very clear manipulation of fairness and unfairness relative to a reference level. When all workers were hired, it was announced that they would be paid 15 euros per hour. There are three treatments. In the fair treatment, the workers are familiarized with their task (entering data into

a library information system), and the subjects were then told that they would be paid 20 euros, not 15, without any explanation for the difference. In the baseline treatment, the subjects were paid 15 euros. In the unfair treatment, the subjects were paid 10 euros, again without an explanation why. Subjects in the fair condition worked about 10 percent more than did subjects in the baseline condition, though the difference is only significant at  $p = 0.165$ . Again, very large individual differences in the baseline output and the small sample size (10 subjects per condition) do not allow for interpreting clear-cut results in the fair treatment. However, in the unfair treatment the reduction in effort relative to the baseline condition is so large (27 percent) that it is significant despite the small sample. These results are in line with evidence from lab studies that find fair treatment has small effects on behavior, but that unfair treatment has large effects on behavior (see, for example, Offerman 2002).

A potential problem in Gneezy and List (2006) and Kube, Maréchal, and Puppe (2008) is that the only manipulation done is by offering higher pay without providing any explanation for the difference. Yet firms go out of their way to stress how well they pay their employees to underscore a salient contrast: comparing the treatment they afford to their employees to what workers would get elsewhere.

Only few studies offer a more specific manipulation of fairness perceptions. Cohn, Goette, and Fehr (2007) implemented a wage increase during a newspaper promotion. The newly launched newspaper hired workers from a marketing agency to distribute their newspaper. In their treatment, the workers were given a CHF 5 increase over their regular pay of CHF 22 and asked to approach the passersby as actively as possible *in return for the higher pay*. In a control treatment, the workers were simply asked to approach the passersby as actively as possible. The promotion was limited in time—in fact, each employee only worked a few days for the newspaper. Thus, the interactions between the workers and the newspaper which implemented the extra pay can essentially be considered a one-shot situation. In an anonymous survey conducted by the marketing agency, the workers clearly stated that they perceived the wage increase as generous, showing that manipulating the fairness perception was effective. The increase in productivity was moderate but statistically significant: depending on the specification,

productivity was around 4 to 5 percent. In a follow-up survey, the authors also asked the workers whether they perceived that the base pay for this job was adequate. Interestingly, the subjects answering that the base pay was inadequate responded significantly more strongly to the intervention. This result is consistent with the prediction from fairness models: the perceptions of the workers who felt that they were treated unfairly were impacted the most by the CHF 5 wage increase. Consequently, they raised their effort more.

A different method of attempting to make the fairness manipulations stronger is reported in Kube, Maréchal, and Puppe (2006). This study varied the form of the gift made to the subjects. The authors hired students to enter data and announced wages would be 36 euros for the three-hour work period. There were three treatments. In the baseline treatment, the students were paid the 36 euros, as announced. In the fair treatment, the subjects were told that they would be paid 43 euros, not 36. In the gift treatment, the subjects were given a Nalgene bottle worth 7 euros at the beginning of the work episode. The idea behind the gift treatment was to manipulate the subjects' perception of kindness: in the case of the Nalgene bottle, it was clear that the experimenter went out of his way to be nice to the subjects by giving them a gift. Thus, if fair and kind treatment increases work effort, this gift treatment should work better than simply paying the subjects more. The results show the usual small effect of a monetary gift: when paid an extra 7 euros, subjects entered approximately 6 percent more data than in the baseline treatment. As usual, the effect is not large enough to be significant. However, in the gift condition, the subjects entered 30 percent more data ( $p < 0.01$ ).<sup>9</sup> A plausible interpretation is that gift exchange in this example is facilitated because it strengthened the signal that the employer cares about the worker, and thus made the difference to the reference transaction more salient. The evidence presented in Kube, Maréchal, and Puppe (2008) suggests that such subtleties are potentially important.

Table 2 summarizes the evidence from field experiments on gift exchange and for each study provides the elasticity of output with respect to the wage. Overall, the field studies on gift exchange show that the hourly wage paid to workers affects their productivity even in short-term jobs that lack the prospect of repeated exchanges—thus indicating that fairness concerns do affect productivity. Moreover, the effect of wage

**Table 2**  
Gift Exchange in One-Shot Interactions: The Field Evidence

Study	Treatment	Elasticity	Type of Task
Kube, M, and P (2008)	Wage Increase	0.30	library task
	Wage Cut	0.82***	library task
Kube, M, and P (2006)	Wage Increase	0.31	library task
	Gift of same value	1.54***	library task
	Gift, value indicated	1.36***	library task
Cohn, G, and F (2007)	Wage Increase	0.16**	newspaper promotion
Gneezy and List (2006)	Wage Increase	0.15	library task
	Wage Increase	0.38	fund-raising task
Bellemare and Shearer (forthcoming)	Wage Increase	0.25***	planting trees
	\$8 to \$16 announced	0.44***	stuffing envelopes
Al-Ubaydli et al. (2008)	Paid \$8 or \$16		

*Source:* Authors' calculations.

*Note:* \*\* and \*\*\* indicate significance at the 5-percent and the 1-percent level, respectively. Elasticities are evaluated as percentage changes relative to the baseline condition.

variations is always in the predicted direction—wage cuts or wage levels that are likely to be interpreted as a violation of a fairness norm cause output reductions (Al-Ubaydli et al. 2008; Kube, Maréchal, and Puppe 2008) while wage variations that may be interpreted as an increase in fairness tend to increase output (Gneezy and List 2006; Kube, Maréchal, and Puppe 2006, 2008). In addition, the empirical pattern supports the view that losses loom larger than do same-sized gains because wage cuts that violate fairness norms trigger stronger output reductions than same-sized wage increases. In Kube, Maréchal, and Puppe (2008), for example, cutting the hourly wage by 5 euros relative to the baseline causes a large output reduction (with an implied elasticity of output with respect to wage of 0.82) while a 5 euro wage increase in the hourly wage leads to much smaller output gains (with an implied elasticity of 0.30). Field experiments that test the effect of wage increases sometimes find positive but insignificant effects (Gneezy and List 2006; Kube, Maréchal, and Puppe 2006, 2008), yet this may also be due to the small number of observations and the large inter-individual performance differences

that are typically found in these studies. However, studies with a larger number of observations (Cohn, Fehr, and Goette 2007) find significant effects. Furthermore, studies that implement gift exchange by providing direct gifts to the workers show surprisingly large positive effects with implied wage elasticities of 1.3 to 1.5 percent.<sup>10</sup> This large effect contrasts sharply with the much lower wage elasticity associated with “simple” wage increases, which are typically between 0.2 and 0.4 percent. A plausible interpretation of these differences is that merely increasing the wage does not automatically trigger an increase in fairness perceptions, while a specific gift is unambiguously associated with the perception of fair treatment. If this interpretation is true it may be possible to magnify the effect of simple wage increases by embedding them in the right context; that is, by making the fairness increase associated with the wage increase more salient to the workers. The relatively high wage elasticity associated with fairness violations suggests that, even in short-term jobs, wage cuts that violate fairness norms may not be profitable. Likewise, the high wage elasticity of direct gift giving suggests that it may be profitable for employers to stimulate the motivation of workers with such gifts, even in one-shot interactions. However, if the low elasticity of simple wage increases should turn out to be the rule rather than the exception, and even if such wage increases will turn out to be generally unprofitable in one-shot interactions, workers’ fairness concerns may nevertheless exert a powerful impact on effort provision and wage setting.<sup>11</sup> The reason is that repeated interactions are a potentially powerful multiplier of the effect of fairness concerns. Whether this conjecture has empirical substance is the topic of the next section.

#### *Gift Exchange in Repeated Interactions*

In the discussion of gift exchange in one-shot interactions in the labor market between firms and workers, we have shown that higher wages have a significantly positive effect on workers’ effort, but we have also emphasized that the positive impact of gift exchange on aggregated market performance is rather limited. However, all the results described so far have abstracted from an important aspect of the labor market: employment relationships are hardly ever spot-market transactions where anonymous trading partners interact only once. Rather, employers and workers have

the option to interact repeatedly with each other. In what follows, we summarize studies which include the long-term nature of relations in labor markets and investigate the interaction of reciprocity and repeated game effects.

#### *Evidence from the Lab*

It can be shown that the presence of a fraction of workers inclined towards fairness allows for the existence of a reputational equilibrium in which not only the fair-minded workers but also the selfish ones are motivated to provide nonminimal effort in many periods of the experimental game. The formal argument for why such an equilibrium can be sustained in a finitely repeated game is related to the result of Kreps et al. (1982). The presence of a fraction of fair-minded workers implies that the firm is willing to pay a worker wages above the reservation level, even in the last period of the interaction—provided there is a sufficiently high belief that the worker is fair-minded. Even though it is certain that all of the selfish types will shirk in the final period, the fair-minded workers will still exert high effort. Therefore, if the probability that the worker is fair-minded is sufficiently high, it pays for the firm to offer high wages to all workers, even in the last period. The prospect of receiving future rents gives selfish workers an incentive to hide their true type from the firm, instead behaving like a fair-minded worker and exerting more effort when offered a high wage. As long as the firm does not detect that a worker is selfish, it will offer him a high wage during every period of the employment relationship, including the last one. By contrast, once a worker reveals that he is selfish, the firm will no longer be willing to pay him more than the reservation wage for exerting minimal effort. This possibility provides selfish workers with a strong incentive to establish the same record, or reputation, as a fair-minded worker.

The first paper that investigates the effect of repeated interactions in a gift exchange setup is Gächter and Falk (2002). They set up a laboratory experiment with two treatments. The baseline treatment involved a sequence of 10 one-shot interactions with a matching scheme that ensured that a particular pair of subjects interacted only once. In the main treatment, each pair of subjects was informed that they would play a 10-times repeated version of this gift exchange game. Thus, in this sec-



ond treatment, each pair of subjects has a common history, and both participants can always condition their actions on their past experience with their partner. If, for example, employers only offer attractive contracts to workers who have always provided high effort in response to high wage offers in the past, then selfish workers have a strong incentive to hide their type and imitate the behavior of fair-minded workers. By providing high effort in response to high wage offers, selfish workers can build up a reputation as fair-minded types. Due to the conditional offering strategy of employers, such a reputation can be of value, as it gives the workers access to profitable future offers from which they would be excluded if their true type were revealed.


The data from this study reveal that repeated game effects are important. In the treatment with repeated interactions the wage-effort relationship is steeper than in the one-shot treatment. As a consequence, average effort levels and market efficiency are significantly higher in the repeated game. Effort levels, which can be chosen between 0.1 and 1, stabilize at about 0.55 and remain there until period nine. In the tenth and final period the effort level drops to approximately the average effort level in the one-shot treatment (0.41). A detailed analysis of individual behavior confirms that this development over time is roughly in line with the reputational explanation put forward above. In both treatments there is a fraction of subjects who are genuinely motivated by fairness concerns. Hence, the repeated game incentives leave intact the subjects' fairness motivation. However, in the repeated game treatment there are also selfish subjects who imitate fair behavior. Thus, the repeated game nature of the treatment disciplines many selfish individuals who would—in the absence of repeated interactions—play uncooperatively. These findings illustrate a fundamentally important point. Although gift exchange alone has only a limited impact on market efficiency, these effects may become larger once firms and workers interact repeatedly with each other.<sup>12</sup>

Brown, Falk, and Fehr (2004) allow long-term employment relationships to arise endogenously in a competitive market environment. In this experiment employers have the opportunity to direct their wage offers to specific workers. They can therefore build up a long-term relationship with a worker by renewing offers to the same worker in consecutive periods. Comparing this treatment to one in which conscious repeated inter-

actions are excluded measures whether the market's participants succeed in endogenously establishing long-term relationships, which serve as an effective effort-enforcement device.

The results of this paper show the importance of endogenous reputation formation in labor markets. Figure 1 displays the distribution of effort choices and average effort over time for both the one-shot condition (OS) and the reputation condition (REP). In the one-shot condition where reputation formation is not possible, the modal effort choice is the minimum level. In contrast, when reputation formation is possible the maximum effort level is most often chosen. Overall, average effort increases from 3.3 in the treatment with one-shot interactions to 6.9 in the treatment with endogenous formation of relationships. The reason for this difference is that in the treatment with fixed identities, many employers succeed in establishing efficient long-term relationships with workers. Employers are mainly interested in interacting with fair-minded workers because these workers are willing to reciprocate high wage offers by providing high labor efforts. Accordingly, most employers are only willing to renew their contract with a worker as long as there is no indication that the worker is selfish. This implies that employers strictly condition the continuation of an attractive position for a worker based on his current effort choice. Since receiving high wage offers generates rents for selfish workers, they are motivated to hide their true type and imitate the fair-minded workers' behavior. In contrast to the situation in one-shot interactions, high wage offers in the relationship condition not only motivate fair-minded workers to provide high effort but also motivate the selfish ones who imitate them. At the end of the experiment, however, the reputation for being fair-minded is no longer valuable for selfish workers, and therefore they no longer hide their type. This leads to a significant drop in performance in the last period.<sup>13</sup>

This study reinforces and extends the findings of Gächter and Falk (2002). In finitely repeated relationships, the presence of a fraction of fair-minded agents who only have a limited impact on performance in one-shot interactions is enough to trigger a strong increase in market performance. The reason is that reputational incentives can motivate selfish agents to imitate the behavior of fair-minded workers. Reputational effects are considerably stronger in the second study, most likely because



**Figure 1**  
Average Effort and Distribution With and Without Reputation Formation

*Source:* Authors' calculations based on Brown, Falk, and Fehr (2004).

*Notes:* The figure displays data from a gift exchange experiment by Brown, Falk, and Fehr (2004). In the one-shot condition (OS) employers cannot distinguish between selfish and fair-minded workers and therefore repeated game effects cannot play a role. In the reputation condition (REP) employers can identify workers such that the endogenous formation of long-term relationships is possible. The top panel of the figure depicts the development of the average effort level over time. The figure's bottom panel shows the distribution of the workers' effort choices over all periods.

of the endogeneity of the long-term relationships. In Gächter and Falk (2002) the participants are forced to interact with the same partner over 10 periods. Thus, while employers can make their offers less attractive within the relationship, they do not have the possibility of terminating the relationship. In contrast, relationships are voluntarily formed in the market of Brown, Falk, and Fehr (2004). Given that there is an excess supply of labor, terminating the employment relationship is a credible threat.

#### *Evidence from the Field*

Field evidence on the role of gift exchange in repeated interactions is more indirect and circumstantial; this is because of the absence of controlled long-term experiments explicitly manipulating wages in a way that creates repeated game incentives.<sup>14</sup> Needless to say, conducting such an experiment would be extremely costly. Perhaps closest to this aim, Bellemare and Shearer (2009) conducted an experiment at a landscaping firm. They implemented a one-time wage gift and examined the workers' response: effort increased ( $p < 0.01$ ). Yet while this experiment taps into a repeated interaction, the treatment is not explicitly geared towards testing how the repeated aspect of the employment relationship affects workers' behavior. The setup is also not ideal, as the landscapers are paid a piece rate; thus increasing their effort also increases the workers' pay.<sup>15</sup>

In the following section we survey evidence from instances in which firms changed the conditions in an ongoing employment relationship. In each case, it is clear that the change negatively affected the workers' fairness perceptions. It is interesting to study these episodes because they can be interpreted as a permanent change in the firm's policy towards its workers and hence the workers' response to this policy change.

There are several striking examples illustrating the potential costs of treating workers in a way that they view as unfair. Krueger and Mas (2004) examine the quality of Bridgestone/Firestone tires manufactured in different plants and years. The particular plant of interest is the one in Decatur, Illinois, which experienced serious labor strife over an extended period of time. The conflict started when the company announced that at all its plants, new hires would be paid less than incumbent employees and that the shift rotations would be altered to a schedule that the work-

ers generally opposed. This announcement triggered a conflict between management and workers at all Bridgeston/Firestone plants. At the Decatur plant, management was particularly aggressive and threatened to hire, and later did hire, replacement workers. This move was seen by the workers as particularly unfair, as it breached the common long-term understanding between management and the workers. Tire manufacturing is still very labor-intensive and depends on high work effort, as the product quality is of utmost importance for the longevity and safety of a tire. The results show clearly that tires manufactured during the Decatur labor strife were of significantly lower quality compared to the same type of tires manufactured at different plants in the same period.<sup>16</sup> A more detailed analysis reveals that an important quality differential was generated immediately after the announcement, even before any of the new policies were put into place. This pattern is particularly supportive of a behavioral model in which workers care about being treated fairly, as the mere intention to act in a way that workers consider unfair triggered the negative response. The data also show that the quality of the tires produced was lowest when many of the union workers had to interact with the nonunionized replacement hires. Thus, it appears that the union workers were the least motivated when they were working side-by-side with employees who accepted the new working conditions. Again, this evidence is supportive of the view that fairness considerations played a key role in understanding the precipitous drop in tire quality at the Decatur plant.

In a similar vein, there is evidence showing that a labor dispute at Caterpillar, a large manufacturer producing construction equipment, tractors, and other vehicles, had a similarly negative impact on production quality (Mas 2007). Negotiations between the union and management broke down after Caterpillar refused to accept a contract that the same union had closed with John Deere, a firm similar to Caterpillar. This move by management was viewed as an attempt to strong-arm the workers into a worse contract and take away rents to which the workers felt entitled. Much as with tire production, a significant share of work on construction equipment is manual, and requires care and effort to produce a high quality product. Mas (2007) shows that, relative to comparable Caterpillar equipment produced outside the United States, the equipment produced in

the United States during the labor strife shows a lower resale value. Mas argues that work effort is an important determinant of quality, and his interpretation is that work effort was lower during the labor strife. Here, again, a conflict erupted between the firm and its workers, and negative consequences followed after Caterpillar announced it would not agree to the new contract. Like the earlier example, this case is consistent with a model in which employees work less hard if they feel treated unfairly.

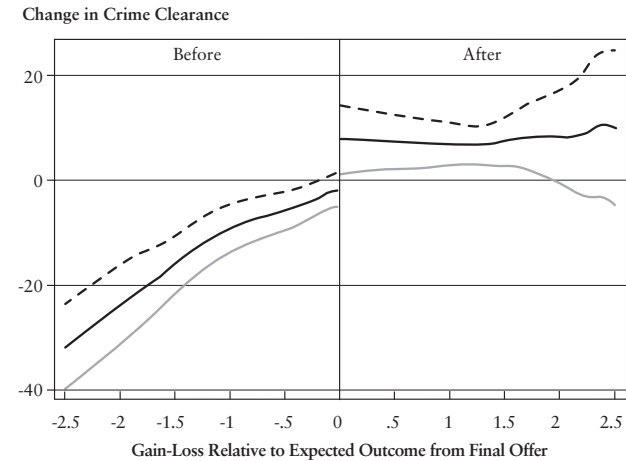
No studies exist examining how the repeated nature of an employment relationship affects the effectiveness of gift exchange in labor markets. However, in a related field, Maréchal and Thöni (2007) conducted an experiment that allowed them to tap into a similar business context. They conducted an experiment in which sales representatives visited stores to sell pharmaceutical products. The treatment in the experiment consisted of a gift—six product samples—that the sales representative gave to the store manager at the start of the visit. Giving the gift strongly increased sales during the representative's visit, and the impact on sales is quantitatively quite large. Average sales per visit are approximately CHF 60 in the baseline condition, while sales in the gift condition are CHF 270. Given that the pharmaceutical firm was willing to visit the stores to realize a sale of CHF 60, giving a gift in this context is highly profitable. Interestingly, the effect is only present if the sales representative had visited the store before. Gifts on initial visits led to no change in sales. This suggests that the gift is tapping into an ongoing relationship between the two agents. As predicted by the theory and by the evidence from lab experiments, this repeated interaction is where the effects of fairness should be largest.

There is also evidence that actions by employers that are considered unfair trigger stronger responses than actions that are considered fair. Mas (2006) examined the outcomes of final-offer arbitration cases that involved a New Jersey police department, in which the union of police officers and the city were unable to negotiate a new contract.<sup>17</sup> In this case, it is rather clear which of the outcomes the workers found most fair as they would not end up in final-offer arbitration if they did not disagree with the offer the employer made. Mas (2006) documented a large and significant decline in many indicators of police performance subsequent to a loss experienced by the police department: the number of crimes cleared decreased significantly, as well as the probability of incarceration

and the sentence length of crimes prosecuted. This is suggestive of an overall decline in police effort in many domains. Mas (2006) calculated the size of a gain or loss relative to the expected outcome of arbitration and plotted the change in effort against this variable. The results are displayed in figure 2. As can be seen, there is a marked drop in clearances if the police lost in final-offer arbitration. It is telling that there is a discrete drop in effort if the police lost, no matter how small the loss is. Further, the figure shows that the decline in effort is highly sensitive to the size of the loss, but not nearly as sensitive to the size of the gain that the police get if the arbitrator rules in their favor. These results are consistent with reference-dependent preferences as discussed earlier.

There are two particular caveats regarding this study's results. The first is the potential in this particular case to underestimate the effect of a gain on effort. It is attractive to study the outcome of arbitration cases because this creates credibly random variation in the terms of the police department's contract. However, most police departments are able to settle on an agreement with their city, and do not end up in final-offer arbitration. It is therefore possible that the only cases that ended up in the sample were those cases in which the police felt strongly entitled to their demands. It is not surprising that in these cases there is only a small positive effect on effort, since the police felt that they simply got what they were entitled to and did not experience a gain from this outcome. The second potential issue is that what gets measured is the consequences of implementing the new contract, not the announcement's immediate effects. In the studies we discussed earlier, many of the negative effects had already materialized when the announcement was made. Such effects are present in the group that ultimately lost and the group that ultimately won the arbitration process and thus are differenced out. Since these effects can be large, the study potentially underestimates the effort reductions caused by treating workers unfairly.

To summarize, this evidence shows that if firms treat workers in ways that are perceived as unfair, this treatment may entail very high costs to the firm. Less is known, however, about the impact of treating workers in a way that is clearly perceived to be fair. While the results in Maréchal and Thöni (2007) are suggestive of positive effects accruing from fair treatment, this remains to be documented in a labor-market setting. In



**Figure 2**  
The Relationship Between Gains and Losses in Wage Bargaining and Effort

Source: Mas (2006).

Note: The horizontal axis measures the gain-loss relative to the expected outcome from final-offer arbitration. The expected outcome is constructed using the predicted probabilities of winning and the two offers from the police union and the city, respectively. The vertical axis shows the change in the number of crime clearances before and after the announcement of the arbitration outcome. The relationship is estimated using local-linear regressions.

particular, it is not clear what role the form of the gift exchange plays: the evidence from one-shot experiments suggests that nonpecuniary gifts work better than simply paying a higher wage. If it is also true that gift exchange in repeated interactions works better when a nonpecuniary gift is used, this may provide a potential explanation for why firms invest so heavily in offering nonpecuniary job benefits. However, more research is needed on this issue. In particular, long-term studies with explicit randomization or credibly exogenous changes in compensation policies, such as adopting a set of policies when a firm is bought by another firm, are needed.

***Internal Labor Markets***

In a pioneering book, Doeringer and Piore (1971) assert that there is a sharp distinction between internal and external labor market arrangements. In particular, workers seem insulated from outside labor-market conditions once they are employed in firms. They argue that these arrangements are difficult to explain from the viewpoint of a neoclassical model: “[W]e doubt that any of the major strands of conventional research will prove capable of assimilating the internal labor market into conventional theory in a useful and meaningful way” (xx). As we argue below, fairness preferences have interesting new implications for how firms set wages over time, giving rise to two of the most important features of internal labor markets.

The evidence on fairness perceptions suggests a shift in what workers feel entitled to as they enter an employment relationship with a firm (Kahneman, Knetsch, and Thaler 1986). While workers who join a firm compared the offer they received from the firm to what they could have gotten otherwise in the labor market when forming their fairness judgments, the evidence strongly suggests that incumbent workers compare any proposed change in the employment relationship to the status quo in order to assess the offer’s fairness.<sup>18</sup> A second important regularity is that there appears to be a strong effect of loss aversion on fairness judgments. For example, a small decrease in the wage does much more damage to fairness judgments than a small increase in the wage does to boost fairness perceptions (Kahneman, Knetsch, and Thaler 1986). It is not clear, a priori, whether loss aversion in fairness judgments applies to the nominal or to the real wage. The survey scenarios in Kahneman, Knetsch, and Thaler (1986) hold the real wage cut constant, showing that over and above the loss in the real wage, individuals consider nominal wage cuts particularly unfair. Shafir, Diamond, and Tversky (1997) also show that nominal wage cuts are perceived as genuinely more unfair. Goette and Huffman (2007b) present evidence that it is the salience of a nominal wage cut—which triggers a strong affective reaction—that informs the fairness judgment. They show that holding the real wage change constant, it is just the wage cuts, not the size of a nominal wage change per se, that influences the affective reaction, a result which is consistent with this interpretation regarding loss aversion.

These features give rise to three specific predictions in the theoretical framework we discussed earlier.<sup>19</sup> The first prediction the model makes is that entry-level wages and the wages of incumbent workers respond differently to changes in labor-market conditions. Entry-level wages should strongly depend on labor-market conditions. If the labor market is tight, workers can find alternative employment at relatively high wages. Thus, a high wage is needed to elicit high effort. When unemployment is high, workers’ outside offers will be worse, and they will be willing to exert high effort for a lower wage. As a consequence, the firm’s optimal entry-level wage is lower when there is slack in the labor market. Conversely, for incumbent workers the reference outcome is the contract that was in place the last period, not the workers’ outside options. This in itself makes the wages of incumbent workers independent of labor-market conditions. The model also predicts cohort effects in wages: because last year’s contract becomes the reference outcome for this year, keeping the same contract is viewed as fair. Thus, if a worker started out with a high entry-level wage, this wage will become the reference wage for the next period, influencing future wage outcomes. The third prediction is related to loss aversion: if workers’ fairness judgments are more strongly affected when they are made worse off, then firms should be reluctant to cut wages. The fairness model is silent as to whether real or nominal wages are the relevant measuring stick for fairness judgments. However the evidence in Kahneman, Knetsch, and Thaler (1986) suggests that it is nominal wage cuts which are considered particularly unfair.

The evidence is generally supportive of the model’s predictions. Several recent studies document that job changers’ wages are more cyclical than job stayers’ wages, including Devereaux (2001); Devereux and Hart (2006); Haefke, Sonntag, and van Rens (2006); and Solon, Barsky, and Parker (1994). In all studies, the wages of individuals entering firms are far more sensitive to business cycle variations. It should be noted that the fairness model does not predict that the incumbents’ wages will never change. In particular, if the firm’s profits rise, so should the incumbents’ wages. Since the studies do not attempt to disentangle shocks that affect the profits of firms (for example, productivity shocks) from other shocks (shocks that only change labor supply), there is no detailed test of this prediction. The study closest to testing this prediction is Beaudry and

DiNardo (1991), who find that current labor-market conditions have almost no effect on current wages, but initial labor-market conditions are a significant determinant of wages. Support also comes from several case studies of firms' personnel files (Baker, Gibbs, and Holmstrom 1994; Eberth, 2003; Treble et al. 2001). Such studies, while less representative, show a much clearer picture of how wages change over the course of a career in ways that are difficult to assess using data from labor force surveys. The evidence of cohort effects is also cleanest in these studies: the picture that emerges is that entry-level wages vary widely between years. Each cohort then gradually increases from the entry-level wage, thus preserving the initial differences in wages.

There is also strong evidence that employers shy away from enacting wage cuts and freeze their employees' wages instead of implementing small wage cuts. Figure 3, using data from Fehr and Goette (2005), shows the distribution of nominal wage changes from two large companies in Switzerland. There are two noteworthy features in the distribution. First, there is a clear drop in the density just around zero. A large fraction of individuals receive a nominal wage change of zero, but almost nobody receives wage cuts. Second, small wage increases are frequent. Hence, there is a clear asymmetry in the distribution of wage changes: wage cuts occur less often than expected, as predicted by the model. The distributions shown here are representative of wage change distributions obtained from personnel files (for example, Altonji and Devereux 2000; Wilson 1999). There are significant measurement problems when moving to more conventional datasets like the Panel Study of Income Dynamics or other labor market surveys. The problem is that wages are typically reported with error (Bound et al. 1994). This problem is accentuated when looking at wage changes and may wrongly lead researchers to conclude that there is a substantial amount of wage flexibility. Indeed, studies that do not control for measurement error find a significant number of wage cuts, though these studies still find a strong asymmetry in the distribution of wage changes (Card and Hyslop 1996; Kahn 1997; McLaughlin 1994). Several methods have been proposed to correct for this problem: some rely on parametric modeling of measurement error (Altonji and Devereux 2000; Fehr and Goette 2005), while others are entirely non-

parametric (Gottschalk 2005). It turns out that the specific form of the correction has very little impact. All studies find, however, that correcting for measurement error is important: once these estimators are applied, the evidence one obtains from the labor force surveys essentially looks like the evidence from personnel files: there are only very few wage cuts.

Alternative models of nominal wage rigidity have been proposed. Malcomson (1997) surveys models of wage bargaining that imply that wages are constant as long as both the firm's and the worker's participation constraint is satisfied. Wages are adjusted only when one of the conditions becomes binding. With positive inflation, this automatically implies that wage cuts will tend to be rare. However, the model also makes the prediction that during deflation, wage cuts would be frequent, and raises rare, with the asymmetry going the other way. However, there is no evidence that wage cuts were frequent in the United States during the Great Depression (Akerlof, Dickens, and Perry 1996).<sup>20</sup>

At the more aggregate level, our framework may also help explain some of the business cycle facts for which the standard model has difficulty accounting. First, the model offers a new source of wage stickiness. For example, the model readily makes the prediction that employment should be more volatile than wages. The reason is that in the face of a positive demand shock, rising employment lowers a firm's average profit (because of diminishing returns to effort). This leads the employees to work harder for a given wage, because their wages are now higher relative to the average profit the firm makes per worker. This increases the workers' effort, but does not require that the firm pay a much higher wage. Therefore, most of the firm's adjustment will come through changes in employment, making the wage relatively unresponsive to changes in demand on the product market (see Danthine and Kurmann 2004). On the other hand, the model also predicts a difference between demand shocks and productivity shocks for wage and employment reactions. In contrast to the demand shock discussed above, a positive productivity shock increases the firm's profit directly. Thus, the workers will lower their effort for a given wage. However, because the workers' effort now becomes more valuable to firms, this reinforces the firm's incentives to raise wages (Benjamin, 2005; Danthine and Kurmann 2004).






Figure 3

The Frequency of Wage Changes for Two Firms

Source: Febr and Goette (2005).

Note: The horizontal axis shows the wage change category in percent. The vertical axis measures the frequency of a particular wage increase.

### *Income Targets and Loss Aversion*

The previous subsection discussed the implications of social preferences and loss aversion in employment relationships characterized by largely noncontractible effort. This forces firms to pay high wages in order to elicit above-minimum effort from workers through the fairness mechanism described earlier. But loss aversion may have implications for the form of optimal contracts even when effort is contractible. The evidence we reviewed on reference-dependent preferences suggests an additional channel that can affect labor supply. Suppose individuals have an income target in mind, and that, as is suggested by the evidence in Heath, Larrick, and Wu (1999), this target inherits all the properties of a reference point. If a worker is paid on a piece rate and works hard enough to surpass the income target, this causes her marginal utility of income to drop discretely because the money now feels like a gain relative to the income target. Now suppose that the piece rate is raised. This makes it easier for the worker to surpass her income target. Hence, the marginal utility of income will, on average, be lower over the day. If this drop in the marginal utility is strong enough, even a purely temporary change in the piece rate can lead to lower effort. This prediction is in stark contrast to the predictions of the standard model in economics. While it is possible that permanent wage increases do not lead to more labor supply because of an income effect (that is, diminishing marginal utility of consumption), this is impossible if the wage increase is temporary, as it only has a negligible income effect in this case.

Armed with these two predictions, Camerer et al. (1997) examined the labor supply of New York City cab drivers as a function of the daily wage. Indeed, they found a very strong negative correlation between implicit hourly wages and hours worked: on “good days,” cab drivers work fewer hours, in line with the prediction from the income-targeting model. Statistically, the effect is highly significant and has been replicated using different samples (Farber 2005) and cab drivers in other countries (for example, Singapore, as studied in Chou 2002). Several possible problems have been raised with this finding. While some have been dealt with adequately (for example, on measurement error, see Camerer et al. 1997; Chou 2002), other problems have remained. One of the trickiest problems that the studies face is that there is no convincing instrument for wages, which has led critics to speculate that supply-side shocks

may drive the variation in wages (Farber 2005). In this case, it would not be surprising to see that cab drivers work shorter hours when wages are high, as it is precisely their desire to work short hours that may have caused the wages to be high.

Other studies have shown that the participation margin of labor supply, meaning the propensity to work at all on a given day, is higher when wages are high (Oettinger 1999). This alone, however, does not invalidate the income-targeting model. The income-targeting model also predicts higher participation when wages are high. The reason for this is straightforward. While the higher wage makes exerting effort on the shift less attractive because the worker finds it easier to surpass her income target, working an extra day has, overall, clearly become more attractive. Therefore, the participation margin of labor supply should increase when wages are high (see Köszegi and Rabin 2006 for a formal treatment of this problem).

Fehr and Goette (2007) conducted a field experiment with bicycle messengers. The messengers were paid a piece rate, and the experiment increased that piece rate by 25 percent during a four-week period. The data from the bike messenger firms allowed Fehr and Goette to examine the overall impact on labor supply, as well as on the participation margin of labor and effort per shift separately. Their results reconcile the earlier findings and provide support for the income-targeting model. They find that the bicycle messengers worked significantly more shifts while they were paid a higher wage. However, they also work less hard while on the shift. In further support of income targeting, Fehr and Goette find that only messengers showing evidence of loss aversion in a separate, unrelated choice experiment reduce effort while receiving the higher wage.

While the previous studies considered changes in the wage or piece rate, Fehr, Goette, and Lienhard (2007) go a different route: they hire temporary workers to enter data into an information system. They manipulate the workers' productivity by slowing down the functioning of the computer interface the workers are using. This causes workers to earn less money than if the slowdown had not occurred. The income-targeting model predicts in this case that employees will work harder because they are farther behind their income target. Indeed, the workers'

effort increases as a function of how long they were delayed. The increase in effort only occurs in the treatment where workers are paid according to the quantity of data entered, not when they are paid fixed wages. This rules out that the delay may have changed the marginal cost of effort. Again, Fehr, Goette, and Lienhard (2007) measure loss aversion in the workers using a simple risky-choice experiment. They find that only the loss-averse workers respond to the slowdown by increasing their effort. For workers behaving in a less loss-averse fashion in the risk experiment, there is no evidence that they work harder subsequently. This evidence is difficult to explain with the neoclassical model, but follows immediately from the income-targeting model.

The income-targeting model also makes the prediction that a windfall gain should change the incentives to work. If by luck a worker earns more than expected, this moves her closer to her target income. Initially, effort should increase, as the marginal utility of income is increasing when below the income target because of diminishing sensitivity. However, having surpassed the target, effort should decrease, as the marginal utility now drops discretely.<sup>21</sup> The standard model, on the other hand, predicts no change in motivation after a windfall gain. Testing the two competing models is difficult because it requires data on effort choices over time. Goette and Huffman (2007a) use data from two bicycle messenger firms in San Francisco that allow them to measure effort over the workday. They use random variation in morning earnings to test whether this affects effort in the afternoon. Goette and Huffman find that a windfall in the morning significantly affects effort in the afternoon. Higher morning earnings lead bike messengers to work harder early in the afternoon, but to work less hard subsequently.

One of the important unanswered questions is why the workers may have a daily, as opposed to a weekly or monthly, income target. Theories are silent on the issue of the choice of the reference frame. A plausible interpretation in the case of the bike messengers and cab drivers is that the income target serves as a rough proxy for the amount of money they need to make per day in order to finance their consumption. In these two applications, the amount made per day is particularly salient to the workers. For example, in the case of the bicycle messengers, they are reminded of how much they have made so far every time they drop off a package

and the customer signs the receipt. This may make daily earnings salient and hence lead to a daily, as opposed to weekly, income target.

In different applications, different reference frames have been proposed. For example, Rizzo and Zeckhauser (2003) examine the labor supply choices of young self-employed doctors. They have data on their income and on what the doctors think is an adequate yearly income, which Rizzo and Zeckhauser argue is the doctors' reference income. They find that if a doctor is below his reference income, he will work more the following year in order, they argue, to close the gap to his reference income. In support of this view, they find that if a doctor is above his reference income, there is no significant change in his work effort over the next year. A plausible interpretation of this result is that the doctors' reference income is derived from their consumption level: they may have set their mind on a certain consumption plan and are willing to exert more effort in order to generate enough income to cling to their reference level of consumption.

This model makes a number of interesting new predictions that can be examined. For example, such a model may make individuals more responsive when piece rates are introduced in an environment in which they cannot perfectly control output. In such cases, large gains in productivity are typically observed (for example, Lazear 2000). Fehr and Lienhard (2007) report evidence from a quasi-experimental change in the compensation scheme involving the removal of a daily guaranteed minimum for one group of workers. The change induced the employees to work much harder, in particular the ones who stood to lose most from the change. The group of workers who initially responded the most readily then showed a gradual decline in effort, while the other group responded less initially and then showed an increase in effort. These results are consistent with a gradual change in reference consumption: the group most affected by the change initially increased labor supply to try to contain the reduction in consumption. This led to a gradual decline in consumption and in the reference point (see Bowman, Minehart, and Rabin 1999 for a fully fledged model). On the other hand, the group least affected gradually increased its effort because the new system generated higher income, thus ratcheting up their reference consumption, which in turn made it optimal for them to increase effort somewhat more.

### 3. Policy Implications

In this section, we discuss several policy implications that are influenced by the behavioral forces that we introduced and for which we examined the evidence in sections 1 and 2. These forces change the way the labor market responds to policy intervention, and we highlight two important areas for policy consideration.

#### *The Importance of Wage Dynamics*

As we argued, the evidence strongly suggests a specific pattern of how workers make fairness judgments in employment relationships. The predictions from the behavioral model for the firms' wage policies are largely confirmed by the data. We discuss two issues that are of clear importance to monetary policy. The first concerns the specific nature of downward nominal wage rigidity and its implications for short-term and long-term tradeoffs between inflation and unemployment. The second highlights a feature of the model that has only recently been studied: there is a new source of persistence that propagates macroeconomic shocks in the economy through the mechanism of internal labor markets.

#### *Downward Wage Rigidity*

The evidence we reviewed when discussing internal labor markets suggests that wages are downwardly rigid. As briefly mentioned above, it is not clear from the evidence on fairness perceptions whether the downward rigidity is in real or nominal wages.

This distinction is very important from a policy perspective, as few studies have assessed the extent of nominal and real wage rigidities. Two related studies, Dickens et al. (2006, 2007) develop a unified model to assess the extent of downward nominal and real wage rigidities (see also Goette, Sunde, and Bauer 2007, which offers a similar approach for Germany, Italy, and the United States). The basic idea behind this approach is to use the features of the wage-change distribution depicted in figure 3; that is, they try to use the drop in the density of wage changes just below nominal zero to assess the extent of downward nominal wage rigidities. Similarly, one can develop an estimator for real wage rigidities. In many countries, such as Britain in the 1980s, there is just as pronounced an

asymmetry around zero *real* wage changes. Such discontinuities near the expected inflation rate are used to estimate the extent of real wage rigidities. The advantage of the work by Dickens et al. (2006, 2007) is that it combines data sources from 13 different countries and uses the same method to estimate downward wage rigidity on all datasets. While there are several studies from different countries, it has been difficult to compare their results, as each study used a different method.

The results show strong evidence of real and nominal wage rigidities in virtually all countries. There are few correlations with institutional variables that predict the type of rigidity. The largest and most robust correlation is with union density: the higher the union density, the more real wage rigidity in an economy, and the less nominal wage rigidity. The United States has much stronger nominal wage rigidity than other countries, particularly those in the euro area.

Wage rigidities are often dismissed as irrelevant given that employment relationships are long run in nature. Therefore, the argument goes, for a given present value of the surplus from an employment relationship, firms can set many different wage paths, including some that have rigid wages, to accommodate the workers' fairness concerns. For example, firms may refrain from enacting a wage cut in one year, but instead not give the worker a wage increase in the next year (see Elsby 2009 for a formal model along these lines). In the context of the model that we discussed, this argument is clearly wrong. The reason is that all the evidence indicates that effort depends on the division of surplus in every period. In the example above, not giving the worker a wage raise in the future would lead the worker to exert less effort. Consequently, not cutting the wage in the current period *does* raise the costs to the firm because offsetting the higher wage in the next period entails costs in the form of lower effort.


Consequently, wage rigidities can have a strong impact on firms' costs and may therefore be effective on the real side of the economy. These rigidities are important for policy for at least two reasons. First, rigid real wages add persistence to monetary shocks in the workhorse model of modern macroeconomics, the New Keynesian framework. In the prototypical model, firms have sticky nominal prices, for example because

of the costs of adjusting prices (for a review, see Goodfriend and King 1997). When setting prices, they take into account the future development of wages. Because these wages determine the firms' marginal costs of production, and because the firms are stuck with the price they set now, the price they set takes into account the expected development of future wages (see Dotsey, King, and Wolman 1999). The evidence on wage rigidities discussed above implies that the real wage is not going to be very responsive to shocks. As a consequence, firms are going to change their price less in response to shocks. But the less they change prices upon the impact of a shock, the stronger and longer are its effects on the real side of the economy. Therefore, downward wage rigidities can contribute to making monetary shocks more persistent. Jeanne (1998) shows that the interaction between price setting and wage rigidities is actually more subtle than the argument given above, and that it takes only a little wage rigidity to make monetary shocks quite persistent, assuming standard degrees of price stickiness. Therefore, using the standard macroeconomic model as the relevant model for policy can be potentially costly, because monetary policy can have effects on the real side of the economy that may be much more persistent than the standard model would predict.<sup>22</sup>

The above argument applies to both real and nominal wage rigidities because the channel through which these affect the real side of the economy results from making marginal costs less responsive in general. However, there is a second, and perhaps more important, channel through which downward *nominal* wage rigidity can affect the real economy. The reason why this may be more important is because downward nominal wage rigidity may affect the long-run unemployment rate, not just the response to shocks. This is because higher wages lead to higher prices charged by firms, depressing aggregate demand and, hence, in equilibrium, employment (Akerlof and Dickens 2007). Akerlof, Dickens, and Perry (1996) build a formal model incorporating such an effect. The empirical estimates of the extent of wage rigidities allow Dickens et al. (2006) to calculate by how much wages have been increased due to wage rigidities. They then estimate a cross-country Phillips curve implied by the model in Akerlof, Dickens, and Perry (1996) that incorporates the

effects of downward nominal wage rigidity. The basic estimate is shown in figure 4. Their results imply that downward wage rigidity substantially increases the long-run unemployment rate. This result is robust to including country-specific intercepts.<sup>23</sup> Further, their results show that nominal and real wage rigidities act in just the same way to increase unemployment. Monetary policy can thus potentially affect long-run output and employment through its impact on how strongly the constraint of downward nominal wage rigidity binds. In particular, tight monetary policy when real wage growth is low could lead to persistent increases in unemployment (also see Fehr and Goette 2005, who find a robust correlation between the impact of wage rigidity on wages and unemployment).

Therefore, the evidence indicates that there is a potential additional constraint on monetary policy: if the labor market is characterized by



**Figure 4:** The Relation between Unemployment and Wage Rigidity  
*Source:* Dickens et al. (2007).

*Notes:* Data is from 13 countries in the International Wage Flexibility Project. The figure shows the relationship between the wage sweep-up, meaning, the extent to which wage rigidities increased the level of wages in a particular year, and unemployment in the same year. Each point is a country-year observation.

strong downward nominal wage rigidity, then keeping inflation low when productivity growth is low may entail a significant employment cost. The evidence also suggests the impact that monetary policy can potentially have on unemployment depends on the structure of the particular labor market: in many countries, particularly in the euro area, there is less evidence of downward nominal wage rigidity than, for example, in the United States or in Switzerland. Rather, real wage rigidities seem to be important. While real wage rigidities still have an adverse effect on unemployment, their effect does not depend on the inflation rate.

Less is known about the shape of wage rigidities when inflation rates are virtually zero over a long period of time. In this case, nominal wage rigidities have been shown to be persistent (Fehr and Goette 2005). There is also some evidence from surveys of inflation expectations that individuals tend to ignore inflation once it becomes low enough (Akerlof, Dickens, and Perry 2000). Indeed, there is also evidence from wage rigidity studies that real wage rigidities tend to become weaker as inflation becomes very low (Bauer et al. 2007). This, in turn, may open the door for yet another channel through which monetary policy may affect the labor market: very low inflation rates may cause individuals to ignore inflation when setting wages, thus giving rise to downward nominal wage rigidity. However, more research is needed to understand how inflation expectations affect wages in a behavioral model of the business cycle.

### *The Consequences of Business Cycles*

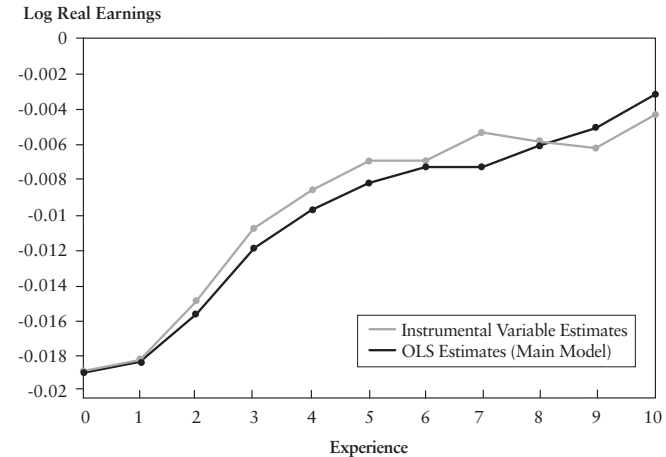
The evidence we reviewed in section 2 shows significant and long-lasting effects of labor-market conditions on individuals' wages. This raises the possibility that even short-run business cycle fluctuations have long-run consequences for the workers in labor market transitions.

This effect may be particularly pronounced for graduating students. Oyer (2008) examines the career choices of Stanford MBA graduates as a function of the stock market, which was highly volatile over the sample period he considered. He finds that MBAs are much more likely to choose employment at an investment bank if the Standard & Poor's 500 index is high than when it is relatively low. Plausibly, investment banks offer more lucrative jobs when business is strong, leading many graduates to take jobs at these firms. The evidence on internal labor

markets then predicts that these employees are able to keep the contract they negotiated—because the firm fears retribution in the form of lower effort if it reneges on the initial contract. Therefore, MBAs entering investment banks in a good year should be more likely to stay in investment banking. Oyer finds that the Standard & Poor’s 500 index has significant effects on job choices in the long run. The level of the Standard & Poor’s 500 in the year of graduation is a highly significant predictor for remaining an investment banker for at least five years.

As we discussed in the section on internal labor markets, entry-level wages are highly volatile over the business cycle. Therefore, business cycles could have long-lasting effects on individuals’ earnings and careers more generally. Oreopoulos, von Wachter, and Heisz (2006) use data on Canadian college graduates to examine the long-run effects of graduating during a recession. They find that the labor-market conditions upon graduation have very strong and long-lasting effects on economic outcomes later. If one graduates in a boom year (with an unemployment rate that is 5 percentage points or lower), initial earnings are about 9 percent higher. Figure 5 displays their results and shows that the effects of unemployment on earnings are long-lasting. After five years, long after the economy has slowed down again, earnings are still 4 percentage points higher and the effect only fades after ten years.<sup>24</sup> The reason why these effects are so long lasting is that initial business cycle conditions change the job-mobility pattern permanently, as one would expect when the firm’s compensation policy is permanently set by initial business cycle conditions.

In summary, the implication of the model that we set forth is that monetary policy may have more persistent effects on the real side of the economy than the standard model would predict. We have offered three channels through which the model outlined above can become relevant in policy considerations: it makes demand shocks more persistent because the behavioral forces discussed above make real wages unresponsive to current economic conditions. This raises the potential of a permanent tradeoff between inflation and employment because the model predicts significant costs to firms from cutting wages. Finally, the model also highlights a new channel through which business cycle fluctuations can be propagated and generate costs for workers over many years.



**Figure 5:** The Impact of Unemployment at Graduation on Log Real Earnings


Source: Oreopoulos, von Wachter, and Heisz (2006).

Notes: The horizontal axis indicates years since graduation. The vertical axis indicates the impact of a 1-percentage point increase in unemployment at graduation in log real earnings. The line labeled “IV estimator” uses freshman year + 4 as the instrument for the graduation year, instead of the actual graduation date in order to avoid endogeneity bias.

#### *Fairness and the Economic Effects of Minimum Wage Laws*

In this subsection we discuss another reason why the psychological forces described in section 2 may be crucial for a better understanding of labor market policy. If people have reference-dependent fairness preferences, policy measures may not only operate by affecting outcomes but also by shifting the relevant reference points. We use the case of minimum wage legislation to illustrate the empirical relevance of this possibility. The minimum wage example is an especially important one because minimum wages are one of the most often-used instruments in labor market policy; for instance, see OECD (1998) for evidence that most labor markets in the developed world are affected by minimum wage laws in one way or another.





**Figure 6**  
The Effect on Wages of Introducing and Removing a Minimum Wage Law  
*Source:* Authors' calculations based on data from Falk, Fehr, and Zehnder (2006).

*Notes:* The figure shows the development of wages over time. Two treatment sequences are displayed. In the first treatment, the introductory sequence (depicted by the solid line) has the participants first experience a situation without a minimum wage, and then the second treatment introduces a minimum wage. In the removal sequence (depicted by the dashed line), the order of the treatments is reversed; the participants first interact in a situation with a minimum wage in the first treatment, and then the minimum wage is removed in the second treatment.

Since minimum wages are so widespread, it is not surprising that for decades economists have been interested in the economic and social consequences of minimum wages. However, despite the remarkable attention the topic has received, at least three frequently reported empirical findings remain puzzling in light of the standard approach in labor economics. First, a number of papers show that minimum wages have so-called spillover effects, meaning that many firms increase wages by an amount exceeding that necessary to comply with the higher minimum wage (see, for example, Card and Krueger 1995; Dolado, Felgueroso, and Jimeno 1997; Katz and Krueger 1992; Teulings 2003; Teulings, von Dieten, and

Vogels 1998). Second, several studies report anomalously low utilization of subminimum wages in situations where firms could actually pay workers less than the minimum wage (see, for example, Freeman, Gray, and Ichniowski 1981; Katz and Krueger 1991, 1992; Manning and Dickens 2002). For example, Katz and Krueger (1991) find that introducing the opportunity to pay subminimum wages to youth has not caused a significant decline in teenage workers' wages. Third, there are several cases in which an increase in minimum wages led to zero or even positive employment effects (see, for example, Card 1992; Card and Krueger 1994; Katz and Krueger 1992; Machin and Manning 1994; OECD 1998; Padilla, Bentolila, and Dolado 1996). This result is surprising because the conventional competitive theory predicts that increases in minimum wages should always reduce employment.

All these effects concern the two most important variables in the minimum wage discussion: wage payments to workers and aggregate employment. Thus, from a policy perspective, a deeper understanding of these puzzling effects of minimum wages would be very desirable.

A recent study by Falk, Fehr, and Zehnder (2006) suggests that the economic consequences of minimum wages can be better understood if the labor market is viewed from the same behavioral perspective that we describe in this paper. However, not only do Falk, Fehr, and Goette (2006) take into account that many workers have reference-dependent fairness preferences, but they also emphasize that labor markets are, in general, not perfectly competitive. This view is based on a recent line of research in labor economics stipulating that imperfect competition is probably the rule rather than the exception in labor markets (see, for instance, Boal and Ransom 1997; Manning 2003). The rationale behind this argument is that labor markets are typically characterized by important frictions (like moving costs, heterogeneous job preferences, or social ties) which prevent the elasticity of an individual firm's labor supply from being close to infinity. Therefore it seems reasonable to assume that firms have at least a certain degree of wage-setting power.

Falk, Fehr, and Zehnder (2006) implement a simple laboratory labor market in which workers' mobility restrictions in combination with heterogeneous fairness preferences give rise to upwardly sloping labor supply schedules at the firm level. They observe that the minimum wage strongly

affects reservation wages, suggesting that it influences what is perceived as a fair wage. After the introduction of the minimum wage there is a strong increase in the reservation wages of subjects playing the role of workers. While almost all reservation wages were clearly below the level of the minimum wage before its introduction, a substantial share of reservation wages are above that level after its introduction. The impact on reservation wages of introducing the minimum wage is in line with the evidence presented in section 2. The mini-ultimatum games of Falk, Fehr, and Fischbacher (2003) revealed that changes in the set of available but unchosen alternatives may have important consequences for the perceived fairness of a specific action. The introduction of a minimum wage takes a whole range of previously possible wage payments out of the strategy set of firms. As a consequence, many subjects seem to perceive a wage payment at the level of the minimum wage—which would have been considered as fair and quite generous before its introduction—as unfairly low after introducing the minimum wage.

The impact of the minimum wage on reservation wages has important implications for the wage-setting strategy of profit-maximizing firms: they are forced to pay wages above the minimum. Thus, the strong impact of the minimum wage on workers' reservation wages provides a possible explanation for the spillover effect empirically observed in field studies. Furthermore, the pattern of reservation wages also shapes the employment effects of the minimum wage. Since firms face upwardly sloping labor supply schedules, they can increase employment if they pay higher wages. However, since the minimum wage not only increases wages but also reservation wages, there is no guarantee that workers will be willing to accept these higher wages. Accordingly, the minimum wage can increase or reduce employment, depending on the relative size of the two counteracting effects. Under the parameters chosen in Falk, Fehr, and Zehnder (2006) the minimum wage has a positive net effect on employment. However, the effect is much smaller than it would have been had workers' reservation wages remained stable.

In contrast to the experimental settings discussed in the previous section, Falk, Fehr, and Zehnder (2006) implemented a labor market with complete employment contracts. However, gift exchange experiments by Brandts and Charness (2004) and Owens and Kagel (2009) show that the

impact of minimum wages on labor supply also prevails if the labor market suffers from contractual incompleteness. Both papers show that the introduction of a minimum wage has two effects. On the one hand, the minimum wage increases average wages, which motivates fair-minded workers to exert more effort. On the other hand, however, the minimum wage also changes the fair-minded workers' willingness to provide effort at a given wage level. It seems that with a law in place that forces employers to pay at least a certain minimum, the same wage is perceived as less fair by the workers than before. As a consequence, the net effect of the minimum wage on effort is ambiguous and depends on the relative size of the two counteracting effects.<sup>25</sup>

In addition, Falk, Fehr, and Zehnder (2006) find that the economic consequences of removing the minimum wage are very asymmetric relative to the effects of its introduction. While workers' reservation wages decrease somewhat after the removal of the minimum wage, they still substantially exceed those before its introduction. It seems that the minimum wage leads to a kind of ratchet effect in workers' perceptions of what constitutes a fair wage. Workers who are used to receiving high wages seem to feel morally entitled to receive them even after minimum wage legislation is abolished. Therefore, the payment of substantially higher wages after the removal of the minimum wage than before its introduction is a profit-maximizing strategy. The asymmetric effect of the minimum wage on reservation wages may explain why firms may find it unprofitable to utilize subminimum wage opportunities—because these opportunities have typically been introduced after a previous increase in the minimum wage.<sup>26</sup>

Of course, laboratory experiments alone will never provide conclusive evidence. However, as the literature on the gift exchange effect shows, effects that have been found in the laboratory may well generalize to field settings outside the laboratory. Thus, if the asymmetric impact of minimum wage laws on reservation wages turns out to be a robust finding, it will have profound consequences. First, it calls into question the basic assumption that labor supply is not affected by the minimum wage. Second, the upward shift in the labor supply curve that is generated by increases in the minimum wage introduces a further potentially employment-limiting aspect of minimum wage increases. Third, the asymmet-

ric impact on reservation wages calls into question the symmetry of the comparative static effects of policy changes. If economic policies generate entitlement effects that respond asymmetrically to the introduction and the removal of the policy, much of what is taught in economic textbooks needs to be rewritten because the introduction of a policy may have effects that prevail even after it is abolished. In the labor market context this means that reductions in the minimum wage are likely to cause much smaller employment effects than one would expect from standard competitive or monopsonistic models.

#### 4. Concluding Remarks

In this study we provide a behavioral view of the labor market. Contrary to standard economic models, our approach accounts for the fact that many employment contracts are incomplete and relational in nature. In addition, we also consider that a substantial fraction of people exhibits reference-dependent fairness preferences shaped by nominal loss aversion. We argue that combining these elements helps us to provide a better understanding of several empirical phenomena which are otherwise considered as puzzles. For example, our approach offers straightforward explanations for well-documented regularities such as downward nominal wage rigidity, the unresponsiveness of incumbents' wages to labor-market conditions, cohort effects, or noncompetitive wage premia. All these phenomena are hard to reconcile with the standard economic model but their explanations arise naturally in our framework.

We also argue that insights from the behavioral economics of the labor market have important policy implications. In many cases, our analysis implies that more responsibility, and power, lies in the hands of policymakers than the standard economic model suggests.

For instance, the behavioral approach to the labor market suggests that downward nominal wage rigidity arises because firms shy away from imposing small nominal wage cuts in response to negative productivity shocks when inflation is low. The reason is that many employees seem to perceive reductions in their nominal pay as very unfair and respond with behavior that is detrimental to the firm. This result implies that monetary

policy can have a permanent effect on the real side of the economy, a finding that contrasts strongly with what a macro model with strictly selfish individuals would predict. Furthermore, downwardly rigid nominal wages also make marginal costs less cyclical and help propagate monetary policy shocks across periods. Thus, the behavioral model provides a simple justification for why marginal costs are not very cyclical, while the standard economic model with selfish preferences must resort to unrealistic assumptions about the elasticity of labor supply or risk-sharing contracts. In addition, the mechanisms of internal labor markets can also propagate shocks across many periods, as has been demonstrated empirically. This channel, too, adds persistence to monetary policy shocks.

In certain cases the behavioral view of the labor market also suggests that policy changes can be hard to reverse. For example, an increase in the minimum wage cannot simply be revoked by subsequently lowering the minimum wage. The reason is that the higher minimum wage leads to a change in what people perceive of as a fair wage. As a consequence, many workers would feel morally entitled to receive a higher wage even if the policy were reversed and therefore would no longer be willing to work for the same wage as before the minimum wage increase. Thus, policymakers need to take into account that today's policies may have important spillover effects on future policy measures.

Overall, we believe that the behavioral approach to the labor market has generated useful insights, and some of these findings have already produced specific policy recommendations—in particular for the cases in which the behavioral model makes qualitatively different predictions than the standard model. It is difficult to quantify many other effects that are important for policy; for example, how internal labor markets affect the persistence of monetary policy. However, as with other recent advances in economics that have policy implications, it takes time to develop models that are specific enough to allow estimation and calibration, but the first steps have already been taken. For instance, Akerlof, Dickens, and Perry (1996) or Danthine and Kurman (2004) offer specific calibrations for policy, and others will follow.

■ *We thank Tyler Williams for excellent research assistance.*

## Notes

1. See Fehr and Tyran (2001); Fehr and Tyran (2008); and Tyran (2007).
2. Most individuals behave in a risk-averse fashion for risks even smaller than in this example. See Holt and Laury (2002) for a recent study.
3. See Rabin (2000) or Fehr and Goette (2007) for details.
4. See Kőszegi and Heidhues (2005) for a formal model of how loss aversion in consumers impacts pricing decisions by firms. See also Rotemberg (this volume) on the implications for price stickiness at the macro level.
5. This does not imply that, in these settings, individuals always behave as the standard model predicts. We return to an important departure from the predictions of the standard model in subsection 3, “Income Targets and Loss Aversion.”
6. Notice that in order for a reputational equilibrium to be sustainable, it is essential that there are fair-minded individuals in the population. So far, we have implicitly assumed that the population fraction of fair-minded individuals is large enough to render a certain degree of gift exchange profitable in one-shot interactions. However, for a reputational equilibrium to be sustainable this assumption is not required. If the fraction of fair-minded individuals is so small that firms would optimally refrain from gift exchange in the one-shot condition, there are reputational equilibria in which some selfish workers start to shirk in the game’s later periods. Since shirking reveals those workers to be the selfish type, they no longer receive wage rents. As a consequence of some selfish workers dropping out, the fraction of fair-minded workers within the group of workers exerting effort increases such that offering wage rents to these workers remains profitable even in the late periods of the game [for details on such equilibria in a related framework (trust game) see, for instance, Camerer and Weigelt (1988) or Brown and Zehnder (2007)].
7. See Benjamin (2005) for a formal model. The following results apply for what he calls perfectly fair workers.
8. Al-Ubaydli et al. (2008) argue that their results should be considered as evidence of gift exchange in the context of repeated interactions because their script told the subjects that it was possible that they may be rehired. However, the fraction of no-shows in the entire study was very high (about 20–30 percent), showing that the workers at the temporary employment agency did not seem to care much about their reputation. This leads us to conclude that incentives arising from repeated interactions are probably not a strong force in this study.
9. One could argue the gift of the bottle causes a positive mood among workers, and that it is because of this mood effect that individuals work harder. However, research shows that positive affective states are not associated with higher productivity (see, for example, Wright and Staw 1999).
10. The elasticity is defined as the percent change in effort in response to a 1-percent wage increase.
11. In order to be profitable in a one-shot setting, the elasticity of output with respect to wages has to be at least 1.0. A complicating feature in all the studies

reported here is that there is no comparable rule by which the wage differences between the treatments are set up. In some studies, researchers may err on the side of caution and implement large wage differences to be sure to find a difference between treatments, whereas the same behavioral response may be obtained with smaller wage increases. Thus, it is difficult to interpret the differences in elasticities between studies. It would be useful to examine the effect of different-sized wage increases within the same framework to get a better sense of what is the profit-maximizing wage.

12. There are other experiments that confirm the role of reputation as an enforcement device. Camerer and Weigelt (1988) study reputation formation in a lending game and Jung, Kagel, and Levin (1994) examine predatory pricing in an experiment where a monopolist faces a series of potential entrants. Both papers find strong evidence for reputation formation in setups with a finite time horizon.

13. The disciplining effect of endogenously formed long-term relationships has also been experimentally investigated in the context of moral hazard in credit markets. See Fehr and Zehnder (2009) and Brown and Zehnder (2007).

14. A large literature, using evidence from case studies and surveys, exists in the organizational behavior literature. Without going into detail, we—again—take issue with many of them for lack of incentives or randomization of the treatments. See Rotemberg (2006) for an excellent review geared towards economists and many references to these studies.

15. As we discuss in subsection 3, “Income Targets and Loss Aversion,” changing the wage on a particular day may have reduced the motivation to exert effort for a different reason.

16. This lower quality translated into many additional tread separations, leading to a large number of deaths and injuries. See Krueger and Mas (2004) for details.

17. In final-offer arbitration, the employer and the employees have to submit a final bid to a third party. This arbitrator then has to pick which of the two bids to implement.

18. A similar effect can be observed in the fairness judgments of price changes. For example, Bolton, Warlop, and Alba (2003) find that, in repeated transactions, the price that a firm charged last was the relevant reference price, much more so than the price the competitors were offering.

19. This section draws heavily on Benjamin (2005), in which proofs of all the statements can be found.

20. Similarly, during a 20 percent deflation between 1879 and 1890 in Switzerland, Imfeld (1991), using evidence from personnel files from five large Swiss firms, reports virtually no wage cuts.

21. This is also known as the goal-gradient effect, as discussed earlier (see Heath, Larrick, and Wu 1999).

22. There is an inherent problem in this class of macroeconomic models that is not solved by simply making the real wage less responsive. The problem is that the data suggest a fair amount of inflation persistence, which is a correlation between current and past inflation, controlling for the driving process of infla-

tion. Fuhrer (2006) makes this argument in detail. See Fuhrer and Moore (1995) for a model giving rise to “true” inherited inflation persistence.

23. The results are remarkably similar to the estimation results in Akerlof, Dickens, and Perry (1996), using a structural model to implicitly estimate the extent of downward wage rigidity from the inflation-unemployment dynamics.

24. One might argue that the timing of graduation is endogenous to the business cycle. However, the results are robust to using the unemployment rate four years after enrollment as an instrument.

25. Owens and Kagel (2009) show that the relative importance of the negative and positive effects of minimum wages on effort strongly depends on the specific experimental setup. They find that the negative effect is more pronounced if the treatments with and without a minimum wage are compared across subjects than when they are compared within subjects.

26. Owens and Kagel (2009) also report findings from sessions where they eliminate a previously introduced minimum wage in their gift exchange setup. However, since the net effect of the introduction of the minimum wage on effort is positive and leads to a Pareto-superior outcome, it is not very surprising that the elimination does not affect outcomes.

## References

- Akerlof, George A. 1982. “Labor Contracts as Partial Gift Exchange.” *Quarterly Journal of Economics* 97(4): 543–569.
- Akerlof, George A. 1984. “Gift Exchange and Efficiency-Wage Theory: Four Views.” *American Economic Review* 74(2): 79–83.
- Akerlof, George A., and William T. Dickens. 2007. “Unfinished Business in the Macroeconomics of Low Inflation: A Tribute to George and Bill by Bill and George.” *Brookings Papers on Economic Activity* 2: 31–47.
- Akerlof, George A., William T. Dickens, and George L. Perry. 1996. “The Macroeconomics of Low Inflation.” *Brookings Papers on Economic Activity* 1: 1–76.
- Akerlof, George A., William T. Dickens, and George L. Perry. 2000. “Near-Rational Wage and Price Setting and the Long-Run Phillips Curve.” *Brookings Papers on Economic Activity* 1: 1–60.
- Akerlof, George A., and Janet L. Yellen. 1988. “Fairness and Unemployment.” *American Economic Review* 78(2): 44–49.
- Akerlof, George A., and Janet L. Yellen. 1990. “The Fair-Wage Effort Hypothesis and Unemployment.” *Quarterly Journal of Economics* 105(2): 255–283.
- Altonji, Joseph G., and Paul J. Devereux. 2000. “The Extent and Consequences of Downward Nominal Wage Rigidity.” *Research in Labor Economics* 19: 383–431.
- Al-Ubaydli, Omar, Steffen Andersen, Uri Gneezy, and John A. List. 2008. “Incentive Schemes to Promote Optimal Work Performance: Evidence from a Multi-Tasking Field Experiment.” Working Paper, University of Chicago.
- Andreoni, James, and John H. Miller. 1993. “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence.” *Economic Journal* 103(418): 570–585.
- Baker, George P. 1992. “Incentive Contracts and Performance Measurement.” *Journal of Political Economy* 100(3): 598–614.
- Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. “The Wage Policy of a Firm.” *Quarterly Journal of Economics* 109(4): 921–955.
- Bauer, Thomas, Holger Bonin, Lorenz Goette, and Uwe Sunde. 2007. “Real and Nominal Wage Rigidities and the Rate of Inflation.” *Economic Journal* 117(524): F508–F529.
- Beaudry, Paul, and John DiNardo. 1991. “The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data.” *Journal of Political Economy* 99(4): 665–688.
- Bellemare, Charles, and Bruce Shearer. 2009. “Gift Giving and Worker Productivity: Evidence from a Firm-level Experiment.” *Games and Economic Behavior* 67(1): 233–244.
- Benjamin, Daniel J. 2005. “A Theory of Fairness in the Labor Market.” Working Manuscript. Available at <http://www.cramton.net/workshop/papers/benjamin-theory-of-fairness-in-labor-markets.pdf>.
- Benjamin, Daniel J. 2006. “The Rotten Firm Theorem.” Working Manuscript. Available at [http://www.gs.cornell.edu/econ/dbenjamin/The\\_Rotten\\_Firm\\_Theorem.pdf](http://www.gs.cornell.edu/econ/dbenjamin/The_Rotten_Firm_Theorem.pdf).
- Bewley, Truman F. 1995. “A Depressed Labor Market as Explained by Participants.” *American Economic Review* 85(2): 250–254. Available at <http://ideas.repec.org/a/aea/aecrev/v85y1995i2p250-54.html>.
- Bewley, Truman F. 1999. *Why Wages Don’t Fall During a Recession*. Cambridge, MA: Harvard University Press.
- Binmore, Ken. 1998. *Just Playing: Game Theory and the Social Contract*, Volume 2. Cambridge, MA: The MIT Press.
- Boal, William M., and Michael R. Ransom. 1997. “Monopsony in the Labor Market.” *Journal of Economic Literature* 35(1): 86–112.
- Bolton, Lisa E., Luk Warlop, and Joseph W. Alba. 2003. “Consumer Perceptions of Price (Un)Fairness.” *Journal of Consumer Research* 29(4): 474–491.
- Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers. 1994. “Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data.” *Journal of Labor Economics* 12(3): 345–368.
- Bowman, David, Deborah Minehart, and Matthew Rabin. 1999. “Loss Aversion in a Consumption-Savings Model.” *Journal of Economic Behavior & Organization* 38(2): 155–178.
- Brandts, Jordi, and Gary Charness. 2004. “Do Labour Market Conditions Affect Gift Exchange? Some Experimental Evidence.” *Economic Journal* 114(497): 684–708.



- Brown, Martin, Armin Falk, and Ernst Fehr. 2004. "Relational Contracts and the Nature of Market Interactions." *Econometrica* 72(3): 747–780.
- Brown, Martin, and Christian Zehnder. 2007. "Credit Reporting, Relationship Banking, and Loan Repayment." *Journal of Money, Credit and Banking* 39(8): 1883–1915.
- Buchan, Nancy, Rachel Croson, and Eric Johnson. 1998. "Country, Culture, and Communication: Extra-Economic Incentives Toward Economic Cooperation." University of Pennsylvania Working Paper No. 98–03–05.
- Buchan, Nancy R., Eric J. Johnson, and Rachel T. A. Croson. 2006. "Let's Get Personal: An International Examination of the Influence of Communication, Culture, and Social Distance on Other Regarding Preferences." *Journal of Economic Behavior & Organization* 60(3): 373–398.
- Bull, Clive. 1987. "The Existence of Self-Enforcing Implicit Contracts." *Quarterly Journal of Economics* 102(1): 147–159.
- Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Quarterly Journal of Economics* 112(2): 407–441.
- Camerer, Colin, and Richard H. Thaler. 1995. "Anomalies: Ultimatums, Dictators and Manners." *Journal of Economic Perspectives* 9(2): 209–219.
- Camerer, Colin, and Keith Weigelt. 1988. "Experimental Tests of a Sequential Reputation Model." *Econometrica* 56(1): 1–36.
- Cameron, Lisa A. 1999. "Raising the Stakes in the Ultimatum Game: Experimental Evidence from Indonesia." *Economic Inquiry* 37(1): 47–59.
- Card, David. 1992. "Do Minimum Wages Reduce Employment? A Case Study of California, 1987–1989." *Industrial and Labor Relations Review* 46(1): 38–54.
- Card, David, and Dean Hyslop. 1996. "Does Inflation 'Grease the Wheels of the Labor Market?'" Working Paper No. 5538. Cambridge, MA: National Bureau of Economic Research.
- Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Unemployment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *American Economic Review* 84(4): 772–793.
- Card, David, and Alan B. Krueger. 1995. *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton, NJ: Princeton University Press.
- Charness, Gary. 2004. "Attribution and Reciprocity in an Experimental Labor Market." *Journal of Labor Economics* 22(3): 665–688.
- Charness, Gary, Guillaume R. Frechette, and John H. Kagel. 2004. "How Robust is Laboratory Gift Exchange?" *Experimental Economics* 7(2): 189–205.
- Chou, Yuan K. 2002. "Testing Alternative Models of Labour Supply: Evidence from Taxi Drivers in Singapore." *Singapore Economic Review* 47(1): 17–47.
- Cohn, Alain, Ernst Fehr, and Lorenz Goette. 2007. "Gift Exchange and Effort: Evidence from a Field Experiment." Manuscript. University of Zurich.
- Danthine, Jean-Pierre, and André Kurmann. 2004. "Fair Wages in a New Keynesian Model of the Business Cycle." *Review of Economic Dynamics* 7(1): 107–142.
- De Quervain, Dominique J.-F., Urs Fischbacher, Valerie Treyer, Melanie Schellhammer, Ulrich Schnyder, Alfred Buck, and Ernst Fehr. 2004. "The Neural Basis of Altruistic Punishment." *Science* 305(5688): 1254–1258.
- Devereux, Paul J. 2001. "The Cyclicity of Real Wages Within Employer-Employee Matches." *Industrial and Labor Relations Review* 54(4): 835–850.
- Devereux, Paul J., and Robert A. Hart. 2006. "Real Wage Cyclicity of Job Stayers, Within-Company Movers, and Between-Company Job Movers." *Industrial and Labor Relations Review* 60(1): 105–119.
- Dickens, William T., Lorenz Goette, Erica L. Groshen, Steinar Holden, Julian Messina, Mark E. Schweitzer, Jarkko Turunen, and Melanie Ward. 2006. "The Interaction of Labor Markets and Inflation: Analysis of Micro Evidence from the International Wage Flexibility Project." *Proceedings*. San Francisco: Federal Reserve Bank of San Francisco. Available at <http://www.frbsf.org/economics/conferences/0603/ilmi.pdf>.
- Dickens, William T., Lorenz Goette, Erica L. Groshen, Steinar Holden, Julian Messina, Mark E. Schweitzer, Jarkko Turunen, and Melanie E. Ward. 2007. "How Wages Change: Micro Evidence from the International Wage Flexibility Project." *Journal of Economic Perspectives* 21(2): 195–214.
- Doeringer, Peter B., and Michael J. Piore. 1971. *Internal Labor Markets and Manpower Analysis*. Lexington, MA: D.C. Heath.
- Dolado, Juan J., Florentino Felgueroso, and Juan F. Jimeno. 1997. "The Effects of Minimum Bargained Wages on Earnings: Evidence from Spain." *European Economic Review* 41(3–5): 713–721.
- Dotsey, Michael, Robert H. King, and Alexander L. Wolman. 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics* 114(2): 655–690.
- Eberth, Barbara. 2003. "Cohort Wage Profiles in an Internal Labour Market." Paper 71. Royal Economic Society Annual Conference. Available at <http://repec.org/res2003/Eberth.pdf>.
- Elsby, Michael W. L. 2009. "Evaluating the Economic Significance of Downward Nominal Wage Rigidity." *Journal of Monetary Economics* 56(2): 154–169.
- Engelmann, Dirk, and Urs Fischbacher. 2009. "Indirect Reciprocity and Strategic Reputation Building in an Experimental Helping Game." *Games and Economic Behavior* 67(2): 399–407.



- Falk, Armin, Ernst Fehr, and Urs Fischbacher. 2003. "On the Nature of Fair Behavior." *Economic Inquiry* 41(1): 20–26.
- Falk, Armin, Ernst Fehr, and Christian Zehnder. 2006. "Fairness Perceptions and Reservation Wages—The Behavioral Effects of Minimum Wage Laws." *Quarterly Journal of Economics* 121(4): 1347–1381.
- Farber, Henry S. 2005. "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers." *Journal of Political Economy* 113(1): 46–82.
- Fehr, Ernst, and Armin Falk. 1999. "Wage Rigidity in a Competitive Incomplete Contract Market." *Journal of Political Economy* 107(1): 106–134.
- Fehr, Ernst, and Urs Fischbacher. 2003. "The Nature of Human Altruism." *Nature* 425 (6960): 785–791.
- Fehr, Ernest, and Simon Gächter. 1998. "How Effective are Trust- and Reciprocity-Based Incentives?" In *Economics, Values, and Organization*, ed. Avner Ben-Ner and Louis Putterman, 337–363. New York: Cambridge University Press.
- Fehr, Ernest, and Simon Gächter. 2000. "Fairness and Retaliation: The Economics of Reciprocity." *Journal of Economic Perspectives* 14(3): 159–181.
- Fehr, Ernst, and Lorenz Goette. 2005. "Robustness and Real Consequences of Nominal Wage Rigidity." *Journal of Monetary Economics* 52(4): 779–804.
- Fehr, Ernst, and Lorenz Goette. 2007. "Do Workers Work More If Wages are High? Evidence from a Randomized Field Experiment." *American Economic Review* 97(1): 298–317.
- Fehr, Ernst, Lorenz Goette, and Marion Lienhard. 2007. "Reference-Dependent Preferences and the Supply of Effort: Evidence from a Field Experiment." Manuscript. University of Zurich.
- Fehr, Ernest, Erich Kirchler, Andreas Weichbold, and Simon Gächter. 1998. "When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets." *Journal of Labor Economics* 16(2): 324–351.
- Fehr, Ernest, George Kirchsteiger, and Arno Riedl. 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics* 108(2): 437–460.
- Fehr, Ernest, and Marion Lienhard. 2007. "The Precautionary Effort Motive: Theory and Evidence." Manuscript. University of Zurich.
- Fehr, Ernst, and Klaus Schmidt. 2003. "Theories of Fairness and Reciprocity: Evidence and Economic Applications." In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*. Volume 1, ed. Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, 208–257. New York: Cambridge University Press.
- Fehr, Ernst, and John-Robert Tyran. 2001. "Does Money Illusion Matter?" *American Economic Review* 91(5): 1239–1262.
- Fehr, Ernst, and John-Robert Tyran. 2008. "Limited Rationality and Strategic Interaction: The Impact of the Strategic Environment on Nominal Inertia." *Econometrica* 76(2): 353–394.
- Fehr, Ernest, and Christian Zehnder. 2009. "Reputation and Credit Market Formation: How Relational Incentives and Legal Contract Enforcement Interact." IZA Discussion Paper Series No. 4351. Bonn: Institute for the Study of Labor. Available at <http://ftp.iza.org/dp4351.pdf>.
- Freeman, Richard B., Wayne B. Gray, and Casey Ichniowski. 1981. "Low-Cost Student Labor: The Use and Effects of the Subminimum Wage Provisions for Full-Time Students." Working Paper No. 765. Cambridge, MA: National Bureau of Economic Research.
- Fuhrer, Jeffrey C. 2006. "Intrinsic and Inherited Inflation Persistence." *International Journal of Central Banking* 2(3): 49–86.
- Fuhrer, Jeff, and George Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110(1): 127–159.
- Gächter, Simon, and Armin Falk. 2002. "Reputation and Reciprocity: Consequences for the Labour Relation." *Scandinavian Journal of Economics* 104(1): 1–27.
- Gächter, Simon, Eric J. Johnson, and Andreas Herrmann. 2007. "Individual-Level Loss Aversion in Riskless and Risky Choice." CeDEx Discussion Paper No. 2007–02. The Center for Decision Research and Experimental Economics, University of Nottingham. Available at <http://www.nottingham.ac.uk/economics/cedex/papers/2007-02.pdf>.
- Gale, John, Kenneth G. Binmore, and Larry Samuelson. 1995. "Learning to be Imperfect: The Ultimatum Game." *Games and Economic Behavior* 8(1): 56–90.
- Genesove, David, and Christopher Mayer. 2001. "Loss Aversion and Seller Behavior: Evidence from the Housing Market." *Quarterly Journal of Economics* 116(4): 1233–1260.
- Gneezy, Uri, and John A. List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments." *Econometrica* 74(5): 1365–1384.
- Goette, Lorenz, and David Huffman. 2007a. "Affect and Cognition as a Source of Motivation: A New Model and Evidence from Natural Experiments." In *Do Emotions Lead to Better or Worse Decisions? A Hedgefoxian Perspective*, ed. Kathleen D. Vohs, Roy F. Baumeister, and George Lowenstein, 267–294. New York: Russell Sage Foundation.
- Goette, Lorenz, and David Huffman. 2007b. "The Impact of Emotions on Wage Setting and Unemployment." In *Do Emotions Lead to Better or Worse Decisions? A Hedgefoxian Perspective*, ed. Kathleen D. Vohs, Roy F. Baumeister, and George Lowenstein, 295–314. New York: Russell Sage Foundation.
- Goette, Lorenz, Uwe Sunde, and Thomas Bauer. 2007. "Wage Rigidity: Measurement, Causes and Consequences." *Economic Journal* 117(524): F469–F477.
- Goodfriend, Marvin, and Robert G. King. 1997. "The New Neoclassical Synthesis and the Role of Monetary Policy." In *NBER Macroeconomics Annual* 12, ed. Ben S. Bernanke and Julio J. Rotemberg, 231–283. Cambridge, MA: The MIT Press. Available at <http://www.nber.org/chapters/c11040.pdf>.

- Goodman, Paul S., and Abraham Friedman. 1971. "An Examination of Adams' Theory of Inequity." *Administrative Science Quarterly* 16(3): 271–288.
- Gottschalk, Peter. 2005. "Downward Nominal-Wage Flexibility: Real or Measurement Error?" *Review of Economics and Statistics* 87(3): 556–568.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze. 1982. "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior & Organization* 3(4): 367–388.
- Haefke, Christian, Marcus Sonntag, and Thijs van Rens. 2006. "Wage Rigidity and Job Creation." Paper 773, 2006 Meeting Papers. Society for Economic Dynamics. Available at <http://repec.org/sed2006/up.26703.1140043700.pdf>.
- Hannan, R. Lynn, John H. Kagel, and Donald V. Moser. 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.
- Hart, Oliver, and Bengt Holmstrom. 1987. "The Theory of Contracts." In *Advances in Economic Theory: Fifth World Congress*, ed. Truman F. Bewley, 71–156. New York: Cambridge University Press.
- Heath, Chip, Richard P. Larrick, and George Wu. 1999. "Goals as Reference Points." *Cognitive Psychology* 38(1): 79–109.
- Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, and Herbert Gintis, eds. 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. New York: Oxford University Press.
- Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath. 2001. "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies." *American Economic Review* 91(2): 73–78.
- Holmstrom, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, & Organization* 7 (Special Issue): 24–52.
- Holt, Charles A., and Susan K. Laury. 2002. "Risk Aversion and Incentive Effects." *American Economic Review* 92(5): 1644–1655.
- Imfeld, D. 1991. "Lohnstarrheit und Soziale Sicherheit des Arbeitsplatzes vor dem Gewerkschaftszeitler: Ausgewählte Schweizerische Arbeitsmärkte im Spiegel Einzelwirtschaf-tlicher Daten aus dem Späten 19. Jahrhundert. Ph.D. Thesis. University of Zurich.
- Jeanne, Olivier. 1998. "Generating Real Persistent Effects of Monetary Shocks: How Much Nominal Rigidity Do We Really Need?" *European Economic Review* 42(6): 1009–1032.
- Jung, Yun Joo, John H. Kagel, and Dan Levin. 1994. "On the Existence of Predatory Pricing: An Experimental Study of Reputation and Entry Deterrence in the Chain-Store Game." *RAND Journal of Economics* 25(1): 72–93.
- Kahn, Shulamit. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review* 87(5): 993–1008.
- Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review* 76(4): 728–741.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem." *Journal of Political Economy* 98(6): 1325–1348.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2): 263–291.
- Katz, Lawrence F., and Alan B. Krueger. 1991. "The Effect of the New Minimum Wage Law in a Low-Wage Labor Market." In *Proceedings of the Forty-Third Annual Meeting*, 254–265. Madison, WI: Industrial Relations Research Association.
- Katz, Lawrence F., and Alan B. Krueger. 1992. "The Effect of the Minimum Wage on the Fast-Food Industry." *Industrial and Labor Relations Review* 46(1): 6–21.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest, and Money*. London: Macmillan.
- Köszegi, Botond, and Paul Heidhues. 2005. "The Impact of Consumer Loss Aversion on Pricing." Discussion Paper No. 4849. London: Centre for Economic Policy Research.
- Köszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121(4): 1133–1165.
- Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson. 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory* 27(2): 245–252.
- Krueger, Alan B., and Alexandre Mas. 2004. "Strikes, Scabs, and Tread Separations: Labor Strife and the Production of Defective Bridgestone/Firestone Tires." *Journal of Political Economy* 112(2): 253–289.
- Kube, Sebastian, Michel André Maréchal, and Clemens Puppe. 2006. "Putting Reciprocity to Work—Positive versus Negative Responses in the Field." Economics Discussion Paper No. 27, University of St. Gallen. Available at [http://www.vwa.unisg.ch/RePEc/uscg/dp2006/DP27\\_Ma.pdf](http://www.vwa.unisg.ch/RePEc/uscg/dp2006/DP27_Ma.pdf).
- Kube, Sebastian, Michel André Maréchal, and Clemens Puppe. 2008. "The Currency of Reciprocity—Gift Exchange in the Workplace." IEW Working Paper No. 377, University of Zurich. Available at <http://www.iew.uzh.ch/wp/iewwp377.pdf>.
- Lazear, Edward P. 2000. "Performance Pay and Productivity." *American Economic Review* 90(5): 1346–1361.
- Levitt, Steven D., and John A. List. 2007. "What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?" *Journal of Economic Perspectives* 21(2): 153–174.

- List, John A. 2006. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy* 114(1): 1–37.
- List, John A., and Todd L. Cherry. 2000. "Learning to Accept in Ultimatum Games: Evidence from an Experimental Design that Generates Low Offers." *Experimental Economics* 3(1): 11–29.
- Machin, Stephen, and Alan Manning. 1994. "The Effects of Minimum Wages on Wage Dispersion and Employment." *Industrial and Labor Relations Review* 47(2): 319–329.
- MacLeod, W. Bentley, and James M. Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57(2): 447–480.
- MacLeod, W. Bentley, and James M. Malcomson. 1993. "Investments, Holdup, and the Form of Market Contracts." *American Economic Review* 83(4): 811–837.
- MacLeod, W. Bentley, and James M. Malcomson. 1998. "Motivation and Markets." *American Economic Review* 88(3): 388–411.
- Malcomson, James M. 1997. "Contracts, Hold-Up, and Labor Markets." *Journal of Economic Literature* 35(4): 1916–1957.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2004. "Disagreement about Inflation Expectations." In *NBER Macroeconomics Annual* 18, ed. Mark Gertler and Kenneth Rogoff, 209–248. Cambridge, MA: The MIT Press.
- Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton, NJ: Princeton University Press.
- Manning, Alan, and Richard Dickens. 2002. "The Impact of the National Minimum Wage on the Wage Distribution, Poverty, and the Gender Pay Gap." London: The Low Pay Commission. Available at <http://www.lowpay.gov.uk/lowpay/research/pdf/amrd.pdf>.
- Maréchal, Michel André, and Christian Thöni. 2007. "Do Managers Reciprocate? Field Experimental Evidence from a Competitive Market." Economics Discussion Paper No. 9, University of St. Gallen. Available at <http://www.vwa.unisg.ch/RePEc/uscg/dp2007/DP-09-Ma.pdf>.
- Mas, Alexandre. 2006. "Pay, Reference Points, and Police Performance." *Quarterly Journal of Economics* 121(3): 783–821.
- Mas, Alexandre. 2007. "Labor Unrest and the Quality of Production: Evidence from the Construction Equipment Resale Market." Working Paper No. 13138. Cambridge, MA: National Bureau of Economic Research.
- McLaughlin, Kenneth J. 1994. "Rigid Wages?" *Journal of Monetary Economics* 34(3): 383–414.
- Odean, Terrance. 1999. "Do Investors Trade Too Much?" *American Economic Review* 89(5): 1279–1298.
- Oettinger, Gerald S. 1999. "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors." *Journal of Political Economy* 107(2): 360–392.
- Offerman, Theo. 2002. "Hurting Hurts More than Helping Helps." *European Economic Review* 46(8): 1423–1437.
- Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2006. "The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates." Working Paper No. 12159. Cambridge, MA: National Bureau of Economic Research.
- Organisation of Economic Co-operation and Development (OECD). 1998. "Making the Most of the Minimum: Statutory Minimum Wages, Employment and Poverty." In *Employment Outlook*, 31–79. Paris: Organisation of Economic Co-operation and Development. Available at <http://www.oecd.org/dataoecd/8/57/2080222.pdf>.
- Owens, Mark F., and John H. Kagel. 2009. "Minimum Wage Restrictions and Employee Effort in Labor Markets with Gift Exchange Present." Manuscript. Available at [http://www.econ.ohio-state.edu/kagel/Owens\\_Kagel\\_MW\\_paper%20revised4\\_30\\_09.pdf](http://www.econ.ohio-state.edu/kagel/Owens_Kagel_MW_paper%20revised4_30_09.pdf). Accessed August 27, 2009.
- Oyer, Paul. 2008. "The Making of an Investment Banker: Macroeconomic Shocks, Career Choice, and Lifetime Income." *Journal of Finance* 63(6): 2601–2628.
- Padilla, A. Jorge, Samuel Bentolila, and Juan J. Dolado. 1996. "Wage Bargaining in Industries with Market Power." *Journal of Economics & Management Strategy* 5(4): 535–564.
- Post, Thierry, Martijn J. van den Assem, Guido Balthussen, and Richard H. Thaler. 2008. "Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show." *American Economic Review* 98(1): 38–71.
- Pritchard, Robert D., Marvin D. Dunnette, and Dale O. Gorgenson. 1972. "Effects of Perceptions on Equity and Inequity on Worker Performance and Satisfaction." *Journal of Applied Psychology* 56(1): 75–94.
- Rabin, Matthew. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica* 68(5): 1281–1292.
- Rabin, Matthew, and Richard H. Thaler. 2001. "Anomalies: Risk Aversion." *Journal of Economic Perspectives* 15(1): 219–232.
- Rizzo, John A., and Richard J. Zeckhauser. 2003. "Reference Incomes, Loss Aversion, and Physician Behavior." *Review of Economics and Statistics* 85(4): 909–922.
- Rotemberg, Julio J. 2006. "Altruism, Reciprocity, and Cooperation in the Workplace." In *Handbook on the Economics of Giving, Reciprocity, and Altruism*, Volume 2, ed. Serge-Christophe Kolm and Jean Mercier Ythier, 1371–1407. Amsterdam: North Holland.
- Roth, Alvin E. 1995. "Introduction to Experimental Economics." In *The Handbook of Experimental Economics*, ed. John H. Kagel and Alvin E. Roth, 3–109. Princeton, NJ: Princeton University Press.

- Roth, Alvin E., and Ido Erev. 1995. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior* 8(1): 164–212.
- Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir. 1991. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *American Economic Review* 81(5): 1068–1095.
- Rubenstein, Ariel. 1979. "Equilibrium in Supergames with the Overtaking Criterion." *Journal of Economic Theory* 21(1): 1–9.
- Seinen, Ingrid, and Arthur Schram. 2006. "Social Status and Group Norms: Indirect Reciprocity in a Repeated Helping Experiment." *European Economic Review* 50(3): 581–602.
- Shafir, Eldar, Peter Diamond, and Amos Tversky. 1997. "Money Illusion." *Quarterly Journal of Economics* 112(2): 341–374.
- Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74(3): 433–444.
- Singer, Tania, Ben Seymour, John P. O'Doherty, Klaas E. Stephan, Raymond J. Dolan, and Chris D. Frith. 2006. "Empathic Neural Responses are Modulated by the Perceived Fairness of Others." *Nature* 439(7075): 466–469.
- Slonim, Robert, and Alvin E. Roth. 1998. "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic." *Econometrica* 66(3): 569–596.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicalities of Real Wages." *Quarterly Journal of Economics* 109(1): 1–25.
- Teulings, Coen N. 2003. "The Contribution of Minimum Wages to Increasing Wage Inequality." *Economic Journal* 113(490): 801–833.
- Teulings, Coen N., Leo von Dieten, and Els Vogels. 1998. *Minimumloon, arbeidsmarkt en inkomensverdeling*. Den Haag: Sdu Uitgevers.
- Treble, John, Edwin van Gameren, Sarah Bridges, and Tim Barmby. 2001. "The Internal Economics of the Firm: Further Evidence from Personal Data." *Labour Economics* 8(5): 531–552.
- Tyran, Jean-Robert. 2007. "Money Illusion and the Market." *Science* 317(5841): 1042–1043.
- Wilson, Beth Anne. 1999. "Wage Rigidity: A Look Inside the Firm." Finance and Economics Discussion Series Paper 22. Washington, DC: Board of Governors of the Federal Reserve System. Available at <http://www.federalreserve.gov/Pubs/feds/1999/199922/199922pap.pdf>.
- Wright, Thomas A., and Barry M. Staw. 1999. "Affect and Favorable Work Outcomes: Two Longitudinal Tests of the Happy-Productive Worker Thesis." *Journal of Organizational Behavior* 20(1): 1–23.

---

## Comments on "The Behavioral Economics of the Labor Market" by Ernst Fehr, Lorenz Goette, and Christian Zehnder

George P. Baker

This paper uses experimental evidence to argue that "fairness preferences" should be built into models of human behavior, and uses this concept to understand several phenomena occurring in labor markets. Fehr, Goette, and Zehnder argue that diverse phenomena such as the form of the wage contract, patterns of wage dynamics, and income targeting can all be understood as manifestations of employees' fairness preferences.

I begin my discussion of this paper by heartily agreeing with the authors' goal of bringing more "behavioral" insights into labor economics. It has long seemed to me that labor economists miss an essential feature of the employment relationship, one that has been known to sociologists and social psychologists for years: people view paid employment as much more than simply an economic transaction. People derive a sense of belonging, identity, and important social benefits from their work and work life. These are notions that remain outside of traditional economic analysis, yet are important to the functioning of organizations and labor markets.

So bringing a more behavioral approach to the functioning of organizations and labor markets will yield, I believe, large benefits. But is adding a preference for fairness the best way to enrich our models? I don't think so. My concern about this approach stems from several sources. First, I think that evidence for a fairness preference rests on shaky foundations. Second, in the context of labor markets and the employment relationship, a repeated game approach to employee and firm behavior seems like a much more sensible way to proceed.

Virtually all of the evidence cited for fairness preferences comes from behavior in one-shot games, especially ultimatum games and gift exchange games. The authors argue that these experiments provide "striking evi-

dence against the self-interest hypothesis” since people in these situations seem to exhibit behavior that helps others, sometimes at their own expense. But is this evidence of a preference for fairness? This implication rests squarely on the assumption that people know, understand, and *feel* that they are truly playing a one-shot game. It is obvious that if an ultimatum game was to be repeated (with the same partner) even one or two times, then the optimal strategy would be for the responder to reject low offers. If your reputation matters at all, then it is not a good idea to be a patsy. This means that rejection of low offers implies fairness preferences (or any other nonself-interested preference) only if the responder has fully internalized the idea of one-shot interactions. The same is clearly true in gift exchange games.

I would suggest that people have a very hard time internalizing the idea of one-shot interactions. People worry about their reputations almost instinctively, and telling them that “there are no long-term consequences” of their play in this game does not make it much easier for them to feel good about playing a strategy (such as being a patsy in the ultimatum game or not reciprocating in a gift exchange) that would be optimal in a one-shot situation. Because of a long evolutionary history of living in small social groups, we have come to care deeply about our reputations and the long-term consequences of our actions: playing a one-shot game is, in a very real sense, an unnatural act. Expecting that laboratory subjects, who are given only a few minutes to understand and think about these games and to fully internalize this unnatural setting, is a tall order. To draw conclusions about their true preferences from their behavior in this unnatural setting is a mistake.

To use evidence from one-shot games to draw implications for human behavior in the employment relationship is especially problematic. If there is any situation in modern life that looks like a repeated game interaction it is employment. People typically interact with their employers over a long period of time, and generally think about the multi-period consequences of their actions at work. The central puzzle of the employment relationship “explained” by Fehr, Goette, and Zehnder—fixed-wage contracts—is more naturally, and parsimoniously, explained by using the tools of repeated game equilibrium.<sup>1</sup>

Fairness preferences are also not needed to explain the well-known patterns of wage dynamics cited by Fehr, Goette, and Zehnder. As discussed in the literature since Doeringer and Piore (1971), a model that includes some firm- and task-specific human capital, on-the-job learning, and worker risk aversion can explain most of the patterns described by Fehr, Goette, and Zehnder (see Harris and Holmstrom 1982; Farber and Gibbons 1996; and Gibbons and Waldman 1999, 2006).

One phenomenon that has been regularly documented in the literature—and not very well explained by standard models—is nominal wage rigidity. It would seem that some behavioral model is needed to explain the sort of money illusion that apparently plagues the employment relationship. Fehr, Goette, and Zehnder argue, disappointingly, that their model is silent on whether nominal or real wages should be downwardly rigid. But I think that this is not so. When using their model to explain business cycle effects, Fehr, Goette, and Zehnder argue that a positive demand shock will not lead to much higher wages: since the firm’s profit per worker does not increase, workers do not feel that working harder for the same wage is unfair. In contrast, when there is a positive productivity shock, the firm’s higher profit level leads fairness-preferring workers to demand higher wages. These conclusions spring from the fact that workers’ fairness preferences are determined by how the *surplus is split* between the firm and the worker.

So the fairness-preference model assumes that workers are sensitive to how a surplus is split. But this assumption implies that workers should be very attuned to inflation: rising prices unaccompanied by rising wages shift the surplus to firms and away from workers. Therefore workers should view inflation unaccompanied by wage increases as unfair and demand higher (nominal) wages. Similarly, in noninflationary times, nominal wage cuts which maintain the same surplus split between firms and workers should not trigger fairness concerns. Thus the fairness-preference model fails to deliver money illusion: on the contrary, it suggests that workers are very attuned to price-level changes. This interpretation is contrary to what we see in the data.

To conclude, I am very sympathetic to the authors’ goal of bringing a more behavioral approach to the study of the employment relationship



and labor markets. But I am not convinced that a model that relies on fairness preferences, derived from observing behavior in one-shot laboratory games, is the right first step.

#### Note

1. The argument that the employment relationship is finite, and therefore any repeated game equilibrium unravels back from retirement, is highly suspect. While the theory of a repeated game equilibrium unraveling is elegant, there is no experimental or field evidence (of which I am aware) of this actually occurring (except perhaps in games played by game theorists!). Consider the implications of this theory in the real world: a 25-year-old employee fails to exert effort at the beginning of his career because he foresees that at as he approaches retirement (decades in the future) there will be an incentive for the firm (whichever firm he works for) to renege on any deal that is made? People do not think this way (or act this way) in the far less socialized environment of the laboratory: the notion that they would do so in the workplace is ridiculous. Furthermore, firms, recognizing the dangers of renegeing on employees as they near retirement, work hard to develop reputations not to renege in this way. Because other employees can observe this renegeing by the firm, this reputational equilibrium is self-enforcing.

#### References

- Doeringer, Peter B., and Michael J. Piore. 1971. *Internal Labor Markets and Manpower Analysis*. Lexington, MA: D.C. Heath and Company.
- Farber, Henry S., and Robert Gibbons. 1996. "Learning and Wage Dynamics." *Quarterly Journal of Economics* 111(4): 1007–1047.
- Gibbons, Robert and Michael Waldman. 1999. "A Theory of Wage and Promotion Dynamics Inside Firms." *Quarterly Journal of Economics* 114(4): 1321–1358.
- Gibbons, Robert, and Michael Waldman. 2006. "Enriching a Theory of Wage and Promotion Dynamics inside Firms." *Journal of Labor Economics* 24(1): 59–108.
- Harris, Milton, and Bengt Holmstrom. 1982. "A Theory of Wage Dynamics." *Review of Economic Studies* 49(3): 315–333.

---

## Comments on "The Behavioral Economics of the Labor Market" by Ernst Fehr, Lorenz Goette, and Christian Zehnder

John A. List

As I sat through the first morning session of the conference, "Implications of Behavioral Economics for Economic Policy," I was struck by the interesting presentations as well as the insightful audience remarks. One particularly astute gentlemen sitting directly at my side made several penetrating comments. One remark related to how one should use experimental methods in economics. He noted that there are distinct advantages in conducting experiments with humans rather than fish, for example, since we can ask humans how they came to make their choice and probe their interpretations of the situation. Of course, this is a valid point. The ichthyologist has no idea what it feels like to be a goldfish when the water temperature suddenly changes from 74 degrees Fahrenheit to 80 degrees Fahrenheit, and the fish might find it difficult to relay that information. Likewise, the chemist has little idea what it is like to transform from uranium-239 into neptunium. Experimental economists are potentially in a more fortunate situation. We are able to study the behavior of people in the laboratory, where we can experience the situations ourselves and ask our experimental subjects about their own experiences.

At the heart of the burgeoning literature that studies fairness in labor markets is experimental evidence of just such a kind. The literature has produced an impressive array of experimental treatments that provide data that have been interpreted as providing strong evidence that many agents behave in a reciprocal manner even when the behavior is costly and yields neither present nor future material rewards (for example, see Fehr, Kirchsteiger, and Riedl 1993; Berg, Dickhaut, and McCabe 1995; Fehr, Gächter, and Kirchsteiger 1997; Fehr and Falk 1999; Fehr and



Gächter 2000; Gächter and Falk 2002; Charness 2004; Brown, Falk, and Fehr 2004; Fehr and List 2004). Such findings have been used to argue, as Fehr and Gächter (2000) put it, that:

Reciprocity has powerful implications for many important economic domains . . . . [T]here are important conditions in which the self-interest theory is unambiguously refuted. For example, in competitive markets with incomplete contracts, the reciprocal types dominate the aggregate results (160).

Using experimental evidence as the cornerstone, the Fehr, Goette, and Zehnder study that I have been asked to comment on prescribes how labor market models should be changed to reflect the manner in which fairness perceptions might affect the labor market. This is an important step in the discovery process, and I laud the authors for their progress. They have made important strides in this study, and in the broader literature with this research agenda. Yet my assigned duty is not to heap praise on these scholars, but to discuss some issues at the heart of this most recent commotion.<sup>1</sup> In this regard, my comment will take a step back and consider more carefully the empirical evidence—both from the lab and field—that has caused this ruckus in the economics community. I should stress at the beginning that I find this research agenda fascinating and that I firmly believe that certain agents have social preferences (see my own field work—for example, List and Lucking-Reiley 2002; List 2004; Landry et al. 2006; Karlan and List 2007).

Nevertheless, the main message of my comment is that the evidence on social preferences from gift exchange games is more mixed than the authors conclude in this study. The gift exchange field studies Fehr, Goette, and Zehnder discuss are confounded in their interpretation—both reputational and social preferences are likely at work. And in those cases when in fact the field data are able to provide a clean measure of social preferences, the effect is found to be small. Lab experiments potentially avoid some of the confounding issues in field studies, but slight perturbations of experimental conditions can dramatically alter behavior, and the important properties of the lab situation are not conducive to fluid generalizations extended to the world beyond the laboratory. In this way, estimating deep preference parameters in the lab and generalizing them

to the field should be done with care, especially in light of the fact that we have no theory to generalize such parameters.

In the end, I view the economics laboratory setting as having a comparative advantage at providing unique qualitative insights. In addition, the lab is able to shed light on what can happen rather than pinpoint what will happen in a certain field situation. For the social preference literature, the lab evidence has certainly highlighted an interesting phenomenon.

### 1. A Framework for Laboratory Experiments in Economics

The basic strategy underlying laboratory experiments in the physical sciences and economics is similar. Yet the fact that humans are the subjects studied in the latter discipline raises fundamental questions about the ability to extrapolate experimental findings beyond the economics lab that do not arise in the physical sciences. Recently with Levitt (2007), I have argued that human decisions are influenced not just by monetary calculations, but also by at least five other factors: 1) the presence of moral and ethical considerations; 2) the nature *and* extent to which one's actions are scrutinized by others; 3) the particular context in which the decision is embedded; 4) the self-selection of the individuals making the decisions; and 5) the stakes of the game.

To make my basic point, I briefly recap the framework we introduced. A utility-maximizing individual  $i$  is faced with a choice regarding a single action  $a \in (0,1)$ . The choice of action affects the agent's utility through two channels. The first effect is on the individual's wealth (denoted  $W_i$ ). The higher the stakes or monetary value of the game, denoted  $v$ , the greater the decision's impact on  $W_i$ . The second effect is the nonpecuniary moral cost or benefit associated with action  $i$ , denoted as  $M_i$ . If, for instance, an individual has strong social preferences, he will derive utility from making charitable contributions.

In practice, many factors influence the moral costs associated with an action, but for modeling purposes we focused on just three aspects of the moral determinant: 1) the greater the negative impact of an action has on others, the more negative the moral payoff  $M_i$ ; 2) the strength of the

social norms or legal rules ( $n$ ) that govern behavior in a particular society and influence behavior; and 3) moral concerns depend on the nature and extent of how an individual's actions are scrutinized. Scrutiny is inherently a multi-dimensional concept, but for simplicity's sake assume that it encompasses only the nature and extent of what is being examined. The nature of scrutiny is exemplified by the presence of an experimenter, who potentially alters the subject's perception of the situation. More broadly, the experimental environment itself might draw upon a different set of expectations than markets. The extent of scrutiny relates to the anonymity of the subject's decision. In the model below, I denote the effect of scrutiny as  $s$ , with higher levels of  $s$  associated with greater moral costs.

Focusing on the case in which utility is additively separable in the moral and wealth arguments, I make the utility function when an individual  $i$  takes action  $a$  as

$$(1) \quad U_i(a, v, n, s) = M_i(a, v, n, s) + W_i(a, v).$$

Solving this simple decision problem yields several predictions, as discussed in Levitt and List (2007). For example, the greater the social norm is against the wealth-maximizing choice or the degree of scrutiny, the larger the deviation from that choice. Furthermore, as the stakes of the game rise, wealth concerns will increase in importance relative to fairness concerns; that is,  $|\partial M/\partial v| < |\partial W/\partial v|$ . Such a framework makes it clear that the greater the extent to which the lab environment mirrors the naturally occurring setting that it is modeling, the more confident one can be that the lab results will be generalizable. If the lab setting diverges from the real-world environment of interest, the model provides a framework for predicting in what direction behavior demonstrated in the the lab will deviate from behavior displayed outside the lab.

## 2. Empirical Evidence

The model can speak to a wide range of experimental results, but its bite is likely to be greatest for those games in which there is the potential for a strong moral component to behavior. Research on social preferences, the

topic of this conference paper, fits the bill. Table 1 highlights a handful of popular empirical approaches—ranging from methods that generate data to techniques used to model data—that have been used to explore preferences. In the leftmost portion of table 1 is laboratory experiments in economics, which are used to generate data; by construction the *ideal* experimental laboratory environment represents the “cleanest test tubes” case. Some might view sterility as a necessary detraction, but sterility serves an important purpose: in an ideal laboratory experiment this very sterility allows an uncompromised glimpse at the effects of exogenous treatments on behavior that takes place in the lab. Of course, making generalizations outside of this domain might prove difficult in some cases, but to obtain the effect of treatment in this particular domain the only assumption necessary is appropriate randomization.

The rightmost part of the empirical spectrum in table 1 includes several examples of empirical models that make necessary identifying assumptions to pinpoint treatment effects from naturally occurring data. These are well-known and need not be further discussed here. Between laboratory experiments and models estimated using naturally occurring data are the various types of economic field experiments that have been introduced recently.<sup>2</sup> As discussed more fully in List (2006a), field experiments represent a useful bridge between the laboratory setting and naturally

**Table 1**  
A Field Experiment Bridge

	Controlled Data		Naturally Occurring Data
	Lab	AFE	FFE, NFE, NE, PSM, IV, STR
■ Lab:	Lab experiment		
■ AFE:	Artefactual field experiment		
■ FFE:	Framed field experiment		
■ NFE:	Natural field experiment		
■ NE:	Natural experiment		
■ PSM:	Propensity score estimation		
■ IV:	Instrumental variables estimation		
■ STR:	Structural modelling		

Source: Author's calculations.

occurring data. Below, we will find that they are crucial when considering generalizability of results.

A useful starting point to summarize the literature is to consider the findings using naturally occurring data, or the rightmost approaches in table 1. An early excellent example is the striking evidence consistent with negative reciprocity on the part of disgruntled Bridgestone/Firestone employees documented by Krueger and Mas (2004). Making use of the variation in product quality induced by the contentious strike and subsequent hiring of replacement workers at Bridgestone/Firestone's Decatur, Illinois plant in the mid-1990s, the authors report that labor strife at the Decatur plant closely coincided with lower product quality. Similarly, Mas (2006) documents persistent adverse effects on police performance following arbitration decisions in favor of the municipality.

The evidence using naturally occurring data is not uniform, however. Chen (2005), who uses a large data set drawn from the Australian Workplace Industrial Relations Survey to explore reciprocity in the workplace, finds little evidence consistent with positive or negative reciprocity. In addition, the empirical results in Lee and Rupp (2006) show that the decreased effort on the part of U.S. commercial airline pilots following pay cuts is a very short-lived behavioral response: even though in the first week after a pay cut frequent and longer flight delays are observed, after the first week there is no difference in airline flight performance. In this manner, these data share important similarities to the short-run effects of shocks observed in the natural field experiment by Gneezy and List (2006).

Moving from the evidence drawn from naturally occurring data to the results from field experiments that explore behavior in repeated play settings, my interpretation of the most recent evidence is that the relationships are *consistent* with models of positive and negative reciprocity (see, for example, Bandiera, Rasul, and Barankay 2005; List 2006b; Cohn, Fehr, and Goette 2007; Marechal and Thöni 2007; Al-Ubaydli et al. 2008; and Bellemare and Shearer 2009). For instance, in List (2006b) I had buying confederates approach dealers on the floor of a sports card show, instructing them to offer different prices in return for sports cards of varying quality. When there was likely to be future interaction and the

consumers could easily certify sports card quality, I found a strong statistical relationship between price and quality provided.

An important consideration, however, is that the theoretical framework proposed above highlights that such data correlations observed in these natural field experiments can operate through social preferences and/or strategic reciprocity. Even if we conclude that the mixed evidence from the naturally occurring data combined with evidence from the natural field experiments painted a picture of an important statistical relationship, by not shutting down the operation of one channel (social preferences or strategic reciprocity), we are confounded in the data interpretation.

Much like we would not ascribe the trendy décor and hip music in Starbucks as evidence of social preferences, the data from the natural field experiments can of course be driven by reputational concerns. Indeed, buttressing this argument is the fact that in my natural field experiment (List 2006b) I also approached the same population of sellers to buy goods that could not be graded. In addition, I approached sellers with whom little future interaction was expected. In both cases a scant statistical relationship between price and quality emerged.

Controlled field experiments as well as laboratory experiments are powerful in the sense that they can *potentially* preclude that reciprocal responses will lead to future material rewards, effectively isolating social preferences. The most common lab game in this spirit is the vast literature on gift exchange, as cited above and reported in Fehr, Gächter, and Kirchsteiger (1997). The experiment is a sequential prisoner dilemma game that has buyers deciding how much money to send to a seller in stage one. In stage two, the seller views this offer (sometimes the offer is multiplied by a factor greater than one) and decides whether to accept it and, if so, what quality to return. The labor-market setting naturally follows if the wage, employer, employee, and work effort are inserted in the relevant portions of the statements. The key behind this approach is that the analyst creates a one-shot environment.

In a natural field experiment testing the gift exchange hypothesis in two actual one-shot labor markets (classifying books in a library and door-to-door soliciting), Gneezy and List (2006) find that worker effort

in the first few hours on the job is considerably higher in a “gift” treatment than in a “nongift” treatment.<sup>3</sup> This result is consonant with the bulk of laboratory evidence on gift exchange cited above. As previously mentioned, however, the result in Gneezy and List (2006) wanes over time and in the long run the gift has a small and statistically insignificant effect. Such insights are in line with results from the psychology literature in that there are important behavioral differences between short-run (*hot*) and long-run (*cold*) decisionmaking (see Loewenstein and Schkade 1999). The notion that positive wage shocks do not invoke long-run effects in effort levels is also consistent with Hennig-Schmidt, Rockenbach, and Sadrieh’s (2006) field experiment (and the lab treatments when employees did not know the surplus division) and Kube, Maréchal, and Puppe (2006). However, it is important to note that the latter do find evidence that negative gifts have short- and long-run effects.

As the preceding discussion suggests, the evidence is generally mixed or negative on the gift exchange relationship in the long run when the repeated game incentives are suppressed. Thus, a useful take-away point is that in one-shot interactions the impact of gift exchange on aggregate market efficiency is small. Equally as important, engaging in gift exchange is unprofitable for principals in these settings. Again, Fehr, Goette, and Zehnder seem to agree, which is an important departure from the earlier literature that argued ferociously about the empirical importance of such preferences in one-shot settings.

### 3. Discussion

Fehr, Goette, and Zehnder stress that the empirical literature shows that “repeated interactions are a potentially powerful multiplier of the effect of fairness concerns (196).” While this claim sounds plausible, I know of no empirical evidence that unambiguously shows this result. More specifically, I am unaware of data that suggests reputational concerns by themselves do not yield the data relationships that are consistent with gift exchange—that is, reputational concerns are able to explain the results without appealing to fairness concerns. Moreover, as alluded to in the empirical data summary, the literature shows that in one-shot environ-

ments the effect of social preferences on aggregate market efficiency is small and trusting actions are unprofitable for the principals.

Yet I am sure that a vector of laboratory design parameters exists that yield a measurable effect of social preferences on market efficiency that is also profitable for the principals.<sup>4</sup> When this set of laboratory results is released, I will interpret the data cautiously because the model in section 1 and the accompanying empirical evidence suggest that a wide range of factors affect the degree to which an agent’s actions will exhibit prosocial tendencies (including the nature and extent of scrutiny), small changes in the way a decision is framed, the stakes involved, self-selection of participants, and artificial restrictions on the action space or duration of play.

List (2006b) presents evidence in favor of some of these conjectures. As briefly mentioned earlier, that study carries out gift exchange experiments in the lab and field that have buyers make price offers to sellers, and in return the sellers select the quality level of the good provided to the buyer. Higher quality goods are costlier for sellers to produce than lower quality goods, but are more highly valued by buyers. In the lab, the results mirrored the typical findings with other subject pools: strong evidence for social preferences was observed.

I then carried out a second lab experiment that maintained the central elements of the gift exchange game, but in a form that was more closely aligned to the context in which sports card trading takes place. The goods exchanged in this lab treatment were actual baseball cards whose market values are heavily influenced by minor differences in condition that are difficult for untrained consumers to detect. If social preferences are present on the part of card sellers, then buyers who offer more money should be rewarded with higher quality cards. When card sellers were brought into the lab to sell their cards, which were subsequently professionally graded, the results paralleled those obtained in the standard gift exchange game with student subjects. However, as noted above, when these same sellers were not aware that their behavior was being scrutinized, the social preferences so routinely observed in the lab were significantly attenuated in the field. The properties of the situation changed in an important manner and this caused sellers to change their behavior in a predictable way.

Other field-generated data yield similar conclusions. For example, making use of personnel data from a leading commercial orchard based in the United Kingdom, Bandiera, Rasul, and Barankay (2005) find that behavior is consistent with a model of social preferences when workers can be monitored, but when workers cannot be monitored, prosocial behaviors disappear. Being monitored proves to be the critical factor influencing behavior in this study. Further, Benz and Meier (2008) combine insights gained from a controlled laboratory experiment and naturally occurring data to compare how individuals behave in donation laboratory experiments and how the same individuals behave in the field. Consistent with the theory in Fehr, Goette, and Zehnder's section 1, they find some evidence of correlation across situations, but find that subjects who in the past have never contributed to charities gave 75 percent of their endowment to the charity in the lab experiment. Similarly, those who never gave to charities subsequent to the lab experiment gave more than 50 percent of their experimental endowment to the charities in the lab experiment.

Gneezy, Haruvy, and Yafe (2004) find that while behavior in a social dilemma game in the laboratory exhibits a considerable level of cooperative behavior, in a framed field experiment that closely resembles the laboratory game they find no evidence of cooperative play—even though both experimental samples are drawn from the same student population. They speculate that unfamiliarity with the task and confusion are two reasons why negative externalities are influential in the lab but not in the field. Such results are consistent with our simple model.

Overall, these results are consistent with the wealth of psychological literature that suggests there is only weak evidence of cross-situational consistency of behavior (see, for example, Mischel 1968; Ross and Nisbett 1991). For instance, Hartshorne and May (1928) discovered that people who cheat in one situation are not the people who cheat in another. If this result spills over to a measurement of prosocial preferences, it means either that (a) there is not a general cross-situational trait called "other regarding," and/or (b) the subjects view one situation as relevant to social preferences and the other as irrelevant. In either case, such insights are consonant with the model, which predicts that factors generating perceptible differences between environments can lead to important behavioral deviations.

#### 4. Concluding Thoughts

Akin to natural scientists, economists have employed experimental methods to lend insights into important phenomena. Recently, the use of laboratory experiments to measure deep preference parameters has grown in popularity, particularly in relation to measuring social preferences. Perhaps the most fundamental question concerning this line of research is whether findings from the lab are likely to provide reliable inferences outside of the laboratory.

In this spirit, the advantage of experimenting with humans becomes a *potentially* serious liability. The choices that individuals make depend not just on financial implications, but also on the nature and degree of others' scrutiny, the particular context in which a decision is embedded, and the manner in which participants are selected to participate. Because the lab systematically differs from most naturally occurring environments on these dimensions, experiments may not always yield results that are readily generalizable.<sup>5</sup>

As I sat down after delivering this message at the conference I was met with a nod, the astute gentleman to my side kindly noted that experimenting with fish does indeed have some advantages after all. "Thank you, Professor Solow" was the only response I could muster.

#### Notes

1. The interested reader should see the older literature as well. Kaufman (1988) provides an edited volume that reviews and assesses the work of four institutional labor economists (John Dunlap, Clark Kerr, Richard Lester, and Lloyd Reynolds). The volume also includes discussions from each of the scholars pertaining to strengths and weaknesses of the literature and the current state of the art. The remarks by Clark Kerr are of particular relevance here. These four economists also wrote a paper titled "Does the New Generation of Labor Economists Know More than the Old Generation?" for Richard Freeman in 1987. Thanks to Alan Krueger for pointing me in this direction and providing the citations. One speculation for why this research agenda is now picking up steam is that what separates the current interest in this topic from the older literature is the recent experimental evidence brought forth.

2. Harrison and List (2004) propose six factors that can be used to determine the field context of an experiment. In doing so, they adopted the term "artefactual" field experiment to denote laboratory experiments with nonstandard sub-

ject pools. Moving closer to naturally occurring data, Harrison and List (2004) denote a “framed field experiment” as the same as an artefactual field experiment but with field context in the commodity, task, stakes, or information set of the subjects. Finally, a “natural field experiment” is the same as a framed field experiment but where the environment is one where the subjects naturally undertake these tasks and where the subjects do not know that they are participants in an experiment. Such an exercise is important in that it represents an approach that combines the most attractive elements of the laboratory setting and naturally-occurring data: randomization and realism.

3. This result is qualitatively similar to other one-shot experiments, but these other studies do not find statistically significant results (see, for example, the field and lab treatments in Kube, Maréchal, and Puppe (2006), and Hennig-Schmidt, Rockebach, and Sadrieh (2006) when employees did not know the surplus division).

4. My intuition is that if one simply increases the multiplier, that will do the trick.

5. This point, of course, applies with equal force to data generated from naturally occurring environments.

## References

- Al-Ubaydli, Omar, Steffen Andersen, Uri Gneezy, and John A. List. 2008. “Incentive Schemes to Promote Optimal Work Performance: Evidence from a Multi-Tasking Field Experiment.” Working Paper, University of Chicago.
- Bandiera, Oriana, Iwan Rasul, and Imran Barankay. 2005. “Social Preferences and the Response to Incentives: Evidence from Personnel Data.” *Quarterly Journal of Economics* 120(3): 917–962.
- Bellemare, Charles, and Bruce Shearer. 2009. “Gift Giving and Worker Productivity: Evidence from a Firm-level Experiment.” *Games and Economic Behavior* 67(1): 233–244.
- Benz, Matthias, and Stephan Meier. 2008. “Do People Behave in Experiments as in the Field?—Evidence from Donations.” *Experimental Economics* 11(3): 268–281.
- Berg, Joyce, John W. Dickhaut, and Kevin A. McCabe. 1995. “Trust, Reciprocity, and Social History.” *Games and Economic Behavior* 10(1): 122–142.
- Brown, Martin, Armin Falk, and Ernst Fehr. 2004. “Relational Contracts and the Nature of Market Interactions.” *Econometrica* 72(3): 747–780.
- Charness, Gary. 2004. “Attribution and Reciprocity in an Experimental Labor Market.” *Journal of Labor Economics* 22(3): 665–688.
- Chen, Paul. 2005. “Reciprocity at the Workplace: Do Fair Wages Lead to Higher Effort, Productivity, and Profitability? Working paper.
- Cohn, A., Ernst Fehr, and Lorenz Goette. 2007. “Gift Exchange and Effort: Evidence from a Field Experiment.” Manuscript. University of Zurich.
- Fehr, Ernst, and Armin Falk. 1999. “Wage Rigidity in a Competitive Incomplete Contract Market.” *Journal of Political Economy* 107(1): 106–134.
- Fehr, Ernest, and Simon Gächter. 2000. “Fairness and Retaliation: The Economics of Reciprocity.” *Journal of Economic Perspectives* 14(3): 159–181.
- Fehr, Ernst; Simon Gächter, and Georg Kirchsteiger. 1997. “Reciprocity as a Contact Enforcement Device: Experimental Evidence.” *Econometrica* 65(4): 833–860.
- Fehr, Ernst, George Kirchsteiger, and Arno Riedl. 1993. “Does Fairness Prevent Market Clearing? An Experimental Investigation.” *Quarterly Journal of Economics* 108(2): 437–459.
- Fehr, Ernst, and John A. List. 2004. “The Hidden Costs and Returns of Incentives—Trust and Trustworthiness among CEOs.” *Journal of the European Economic Association* 2(5): 743–771.
- Freeman, Richard. 1987. “The Postwar View of Labor Markets and Wage Determination.” In *How Labor Markets Work: Reflections on Theory and Practice* by John Dunlop, Clark Kerr, Richard Lester, and Lloyd Reynolds, ed. Bruce E. Kaufman, 205–232. Lexington, MA: D.C. Heath.
- Gächter, Simon, and Armin Falk. 2002. “Reputation and Reciprocity: Consequences for the Labour Relation.” *Scandinavian Journal of Economics* 104(1): 1–27.
- Gneezy, Uri, and John A. List. 2006. “Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments.” *Econometrica* 74(5): 1365–1384.
- Gneezy, Uri, Ernan Haruvy, and H. Yafe. 2004. “The Inefficiency of Splitting the Bill: A Lesson in Institution Design.” *Economic Journal* 114(495): 265–280.
- Harrison, Glenn W. and John A. List. 2004. “Field Experiments.” *Journal of Economic Literature* 42(4): 1009–1055.
- Hartshorne, Hugh, and Mark A. May. 1928. *Studies in Deceit*. New York: Macmillan.
- Hennig-Schmidt, Heike, Bettina Rockenbach, and Abdolkarim Sadrieh. 2005. “In Search of Workers’ Real Effort Reciprocity – A Field and a Laboratory Experiment.” GESY Discussion Paper No. 55. Department of Economics, University of Mannheim. Available at <http://www.sfbtr15.de/dipa/55.pdf>.
- Karlan, Dean and John A. List. 2007. “Does Price Matter in Charitable Giving? Evidence from a Large-Scale Natural Field Experiment.” *American Economic Review* 97(5): 1774–1793.
- Kaufman, Bruce. 1988. “The Postwar View of Labor Markets and Wage Determination.” In *How Labor Markets Work: Reflections on Theory and Practice* by John Dunlop, Clark Kerr, Richard Lester, and Lloyd Reynolds, ed. Bruce E. Kaufman, 145–203. Lexington, MA: Lexington Books.



- Krueger, Alan B., and Alexandre Mas. 2004. "Strikes, Scabs, and Tread Separations: Labor Strife and the Production of Defective Bridgestone/Firestone Tires." *Journal of Political Economy*, 112(2): 253–289.
- Kube, Sebastian, Michel André Maréchal, and Clemens Puppe. 2006. "Putting Reciprocity to Work—Positive versus Negative Responses in the Field." Economics Discussion Paper No. 27, University of St. Gallen. Available at [http://www.vwa.unisg.ch/RePEc/uscg/dp2006/DP27\\_Ma.pdf](http://www.vwa.unisg.ch/RePEc/uscg/dp2006/DP27_Ma.pdf).
- Landry, Craig, Andreas Lange, John A. List, Michael K. Price, and Nicholas Rupp. 2006. "Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment," *Quarterly Journal of Economics* 121(2): 747–782.
- Lee, Darin, and Nicholas G. Rupp. 2007. "Retracting A Gift: How Does Employee Effort Respond To Wage Reductions?" *Journal of Labor Economics* 25(4): 725–762.
- Levitt, Steven D., and John A. List. 2007. "What do Laboratory Experiments Measuring Social Preferences Tell Us about the Real World?" *Journal of Economic Perspectives* 21(2): 153–174.
- List, John A. 2004. "Young, Selfish, and Male: Field Evidence of Social Preferences." *Economic Journal* 114(492): 121–149.
- List, John A. 2006a. "Field Experiments: A Bridge between Lab and Naturally Occurring Data." *Advances in Economic Analysis & Policy* 6(2), Article 8.
- List, John A. 2006b. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy* 114(1): 1–37.
- List, John A., and David Lucking-Reiley. 2002. "Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign." *Journal of Political Economy* 110(1): 215–233.
- Loewenstein, George. 2005. "Hot-Cold Empathy Gaps and Medical Decision Making." *Health Psychology* 24(4, Suppl.): S49–S56. Available at <http://www.rci.rutgers.edu/~gbc/psycDM/Loewenstein2005.pdf>.
- Loewenstein, George, and David Schkade. 1999. "Wouldn't It Be Nice? Predicting Future Feelings." In *Well-Being: The Foundations of Hedonic Psychology*, ed. Daniel Kahneman, Ed Diener, and Norbert Schwartz, 85–105. New York: Russell Sage Foundation.
- Maréchal, Michel André, and Christian Thöni. 2007. "Do Managers Reciprocate? Field Experimental Evidence from a Competitive Market." Economics Discussion Paper No. 9, University of St. Gallen. Available at <http://www.vwa.unisg.ch/RePEc/uscg/dp2007/DP-09-Ma.pdf>.
- Mas, Alexandre. 2006. "Pay, Reference Points, and Police Performance." *Quarterly Journal of Economics* 121(3): 783–821.
- Mischel, Walter. 1968. *Personality and Assessment*. New York; Wiley.
- Ross, Lee and Richard E. Nisbett. 1991. *The Person and the Situation; Perspectives of Social Psychology*. New York: McGraw-Hill.

## 5

---

## Behavioral Economics and the Housing Market