Wage Policies, Incentive Schemes, and Motivation

Gary Charness, Michael Cooper, and J. Lucas Reddinger*

7 January 2020

Abstract:

This chapter reviews the literature on both financial incentives and non-monetary motivations (social preferences) that affect effort and performance in the labor market. Neo-classical theory implies that higher levels of financial incentives should monotonically increase effort. While this is true over much of the wage range, experimental and applied evidence shows that both very low and very high levels of incentives can violate this theory. Low levels of pay may crowd out intrinsic motivation, while high levels of pay may cause involuntary arousal, or "choking", that reduces performance; however, recent literature has sometimes found contradictory evidence on these effects. Social preferences for fairness and reciprocity also impact the effort decisions of workers. With respect to wage changes, the literature shows that negative reciprocity is typically considerably stronger than positive reciprocity. Additionally, perceptions of fairness in terms of co-worker wages can impact effort. Taken as a whole, the behavioral effects of financial and non-financial motivators depend heavily on the context of the labor market and the types of workers and tasks involved.

* Charness, University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106-9210, charness@econ.ucsb.edu. Cooper, University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106-9210, michael_cooper@ucsb.edu. Reddinger, University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106-9210, reddinger@ucsb.edu.

1. Introduction

What are the motivations that underlie effort and performance in the labor market and elsewhere? This chapter will discuss both financial and social-preference motivations as factors that affect behavior, considering also worker responses in experimental labor environments. The review will consist primarily of experimental papers and will depart from the standard neo-classical viewpoint in that it will be clear that financial factors alone cannot explain the observed behavior.

The standard view in economics is that stronger financial incentives should induce higher levels of effort, and that people should only be concerned with maximizing their own monetary resources. However, experimental evidence on financial incentives and the prevalence of charitable contributions have long called these assumptions into question. This chapter will begin with a review of the effects of financial incentives on effort, focusing on the unusual behavioral impacts of very low and very high levels of pay. There is experimental evidence that paying small amounts for performance may be worse than paying nothing at all and relying on an individual's sense of honor or responsibility, although there is dissenting evidence from other studies. There is also evidence that very high levels of performance-based pay can adversely affect performance, perhaps representing the effect of the extreme pressure that could be present. However, the recent literature suggests that the context of the labor market impacts whether these effects occur.

The chapter then proceeds to discuss non-financial motivators for effort and performance. The social-preference literature discusses when and how people will trade off their own material wealth to help or hurt others. This is a vast literature and this review can only provide a snapshot of the theory and the applications. In addition to discussing the role of social preferences in worker-firm relationships, this chapter will also review unusual behavioral impacts. The chapter then closes with a more general discussion of worker responses to various wage circumstances, including wage cuts, wage raises, and wage inequality between workers. Note that while identity and group membership are important considerations for worker productivity, these topics are omitted here; see instead the chapter by Li titled "Group identity, in-group favoritism, and discrimination" and also Charness and Chen (2019).

2. Financial incentives and performance

Economists generally expect that increasing financial incentives for performance should increase effort and productivity. This is a central assumption in the neoclassical principal-agent model, which forms the basis for analyzing effort in personnel economics. This model assumes that workers have increasing and convex costs of effort, and that they set the marginal cost of effort equal to the marginal benefit (i.e., the piece rate) to determine the optimal level of effort. This implies that increasing the piece rate monotonically increases effort. Indeed, the solution to the traditional model indicates that the principal should sell the job to the worker, thereby providing a 100% piece rate and leading to the efficient level of effort (Jensen and Meckling 1976).

A large body of theoretical and applied work has emerged out of this basic principal-agent model to analyze the impact of various types of incentives on effort. Indeed, many forms of compensation contracts are found in practice, including fixed wages per hour, salaries, performance-based pay and bonuses, team-based pay, and tournaments. The scope of this chapter will be limited primarily to relationships between a single worker and a firm, and thus will not cover in detail team-based or tournament-based pay schemes which must consider the effort and pay of multiple workers at once.

Applied work has generally confirmed that in practice, productivity increases when firms move from fixed payments to performance-based pay. Perhaps most famously, Lazear (2000) analyzes administrative data from the Safelite Glass windshield repair company during their transition from hourly wages to piece rates. After controlling for various observable characteristics including seasonal effects and long-term trends in productivity, he finds that output increases by about 44% when the firm moves to performance-based pay.

Similar results supporting the link between incentives and effort have been found in experimental environments. For example, Dickinson (1999) runs a lab experiment in which subjects are paid piece rates to transcribe paragraphs of text. Piece rates are varied both up and down over the course of multiple work sessions, and subjects on average show improvements in productivity when piece rates are increased. Dohmen and Falk (2011) also find that offering performance pay increases output, but do so in an environment with self-sorting of subjects into fixed pay or piece rates. In addition to the pure incentive effect (a higher marginal benefit of

effort leads to higher effort), the authors suggest that the performance differences are strongly driven by sorting effects: higher productivity subjects sort into variable pay schemes rather than fixed pay schemes. This echoes the result of Lazear (2000), who finds that roughly half of the 44% productivity increase at Safelite under piece rates could be attributed to attracting and retaining workers with higher productivity.

These papers among many others have solidified the received wisdom that in general, effort and performance typically increase with financial rewards. However, there are certain contexts in which it seems that increasing incentive pay does not monotonically increase effort. In some cases, there may be a negative effect from increased financial incentives, particularly at the extremes of the incentive-pay distribution. Really small performance incentives can elicit lower effort than no incentive pay at all, and very high performance pay can also be ineffective. In the first case, the thrust of the story is that when people are motivated by intrinsic motivation to perform a task, providing extrinsic motivation can change the perception of the activity and thereby crowd out the original intrinsic motivation. In the second case, people may choke under pressure or even have a backward-bending labor supply curve at very high rates of compensation. People are motivated by many factors besides financial incentives, and these alternative motivations can interact with the level of pay generosity in ways not predicted by standard economic theory.

2.1. Paying too little

A primary way in which low incentive pay can elicit worse effort than no pay is by crowding out intrinsic motivation. Titmuss (1970) made this point in the context of blood donations, arguing that paying blood donors would reduce the amount and quality of donations. He argued that individuals are already intrinsically motivated to donate, meaning that they do the task purely for the internal satisfaction, rather than for any external reward.

A strand of psychology literature emerged exploring this crowding-out effect, which generally found that monetary incentives reduced intrinsic motivation to perform a task (see Deci et al. 1999 for a meta-analytic review of early research on this topic). The outcomes of interest to the psychologists shaping this literature were quite different from those of the economics literature that would later emerge: these psychology studies typically directly measured intrinsic motivation through survey questions or measured how long subjects continued voluntarily to

perform a task after monetary incentives were removed. In contrast, economists became more interested in whether providing small monetary incentives during an intrinsically-rewarding task could reduce effort or output compared to no monetary incentives.

The insight for a financial incentive changing the perceived nature of an activity can be seen in these paired examples taken directly from Gneezy (2003):

Scenario 1: You live in Chicago where recycling soda cans is not extrinsically rewarded. On a freezing morning, you see your friend with a large bag full of soda cans on her way to the recycling container. Clearly you are impressed by her devotion to the environment.

Scenario 2: The same as above, but a five cent reward for each recycled soda can is in place. Now you see your friend carrying a large bag of soda cans to the recycling container (where she can collect her reward).

In the first case, you are likely to consider that your friend is concerned about the environment, while in the second case, it seems that your friend is simply a tightwad. A second pair of examples also taken directly from Gneezy (2003) shows how bringing financial incentives into the picture totally changes the feel of the situation by switching domains:

Scenario 3: Your child is in a daycare center and you are supposed to pick her up at 4:00 p.m. There is no specified sanction for picking up the child late.

Scenario 4: The same as above, but a $3 fine is used as a sanction for late-coming parents

Clearly these environments give different signals. One is on one's honor in the first case, while in the second case the small fine suggests that this is an unimportant violation. One may very well value one's honor more highly than $3. In fact, this was also the case in an experimental study. Small monetary sanctions led to more violations than were present with no monetary sanctions. Nevertheless, the evidence on crowding-out is somewhat mixed, as seen below.

Although this review will focus primarily on applied evidence, some economic theorists have provided valuable insights about wages crowding out intrinsic motivation. Benabou and Tirole (2003) construct a model to explain crowding out – their intuition relies on asymmetric information, so that a principal providing incentives for an agent is a signal about the unattractiveness of the task or the agent's inability to complete the task easily. Frey (1997) uses a similar principal-agent model to propose that increased incentive pay can reduce effort through a

crowding-out effect, while Benabou and Tirole (2006) show how financial incentives can similarly crowd out pro-social behavior by reducing the reputational gains from altruistic actions.

Crowding out may also occur due to a non-financial choice of employers: the level of monitoring and control to exert over worker effort. Although principal-agent theory suggests that monitoring worker output will increase effort, in some contexts this sort of control may crowd out intrinsic motivation and reduce effort. Frey (1993) provides both theoretical and empirical backing to this effect. He argues that when employers exert what is perceived as excessive control over workers, it may violate an implicit contract as understood by the workers. This signal of distrust may lead to lower effort. Such crowding out, Frey suggests, is most likely to occur when the work relationship is perceived as personal rather than abstract or competitive.

Reduced worker effort due to monitoring and control has been found in lab environments. Falk and Kosfeld (2006) set up a principal-agent game in which principals have the choice to restrict the agents' minimum effort, and agents decide on a level of effort (with a monetary cost of effort rather than a real effort task). In this study, agents tend to decrease effort when principals choose to impose a minimum level of effort. Follow-up surveys indicated that controlled agents perceived it as a signal of distrust and a limitation on their autonomy. Dickinson and Villeval (2008) add nuance to this discussion with an experimental design that includes both "distant" and "interpersonal" treatments. They find that on average, monitoring does increase worker effort as predicted by the standard principal-agent model. However, crowding out of effort did occur when principals exerted high levels of monitoring, and this was more likely to occur in the interpersonal treatment where subjects interacted face to face for five minutes before the main task.

Non-financial motivators such as control and monitoring clearly may reduce effort in ways not predicted by the standard model. However, this effect depends on the context of the work relationship and is sensitive to the precise level of control exerted and the perceptions of workers. The remainder of this section will focus on the effects of purely financial motivators on effort.

2.1.1. Positive evidence on crowding-out

Gneezy and Rustichini (2000a) conducted the seminal economics experiment on the detrimental effects of low pay, reported in the aptly-titled paper, "Pay enough or don't pay at

all." The paper was motivated by the contradiction between economic models of an upward sloping labor supply and the results from experimental psychology on how payments can reduce self-reported motivation and voluntary effort levels. They run two separate experiments: one on high-school students soliciting donations for charity door to door, and another on university students completing an IQ test in the lab. In both experiments, they find that the donations collected or the number of questions correct decreases in treatments with very low incentive pay compared to no incentive pay.

Gneezy and Rustichini (2000b) consider the daycare scenario mentioned above. Initially there was no prescribed sanction for being late, but one knew that this might force daycare employees to stay after closing time. A $3 fine for being late was then introduced, leading to an increase in the likelihood that parents would be late. Perhaps even more concerning is the fact that this drop in punctuality did not reverse when the fine was eventually removed. It appears that parents interpreted the small fine as meaning that being late was not a serious issue and that one could assuage any guilt by simply paying this fine.

Gneezy (2003) follows up on these studies and supports the principle that small incentives are less effective than no incentives even in the negative domain, as mentioned in the preceding paragraph. He considers a game in which a first party could choose to send some amount of her endowment to a second party. There were five treatments that varied the possible response. In the control treatment, no response was possible. In the low-reward treatment, a one-unit expenditure by the responder led to an increase of 1.5 units for the first mover; in the high-reward treatment, the increase was instead 5 units. Conditions were exactly the same in the two punishment treatments, except that the expenditure decreased the first-mover's money by either 1.5 or 5 times the amount. The intuition is that monetizing per se changes the perception of the nature of the environment, but that making the sanction quite large will nevertheless more than overcome this impediment. The results support this intuition. When responders could implement a low fine or a low reward, the average transfers from first movers were 6.9 and 7.2, respectively, compared with transfers of 9 when no response was possible. However, when high fines or high rewards were available to responders, the first movers transferred more than in the dictator setting (an average transfer of 10.4 with a high fine, and 12.7 with a high reward). This suggests that while small sanctions and rewards are less effective than no incentives at all, once

these incentives are increased in strength they are indeed effective financial motivators.

Other experiments have followed up on these results and found that low levels of incentive pay can reduce effort compared to no pay or a fixed level of pay. These experiments have uncovered interesting insights by varying the framing of payments, the framing of the work, the kind of task, and other variables. A common theme, although not universally supported, is that low levels of pay reduce effort primarily when the task is already intrinsically motivating – it's interesting, fun, or inherently rewarding. Additionally, low levels of pay are less likely to reduce effort when payments are framed as reciprocal social gifts rather than cash payments for labor.

Ariely and Heyman (2004) vary the framing of payments as cash or gifts to determine in which case crowding-out is more likely. Subjects perform a computerized real-effort task for varying levels of cash or jellybeans. They found that a low monetary payment resulted in less effort than no payment at all, but there were no significant differences in effort between payment levels when payments were in the form of jellybeans. (An important caveat is that the payments occurred before the task and were not contingent on the precise level of effort, so they were not technically piece rates.) Yam (2013) finds a similar result on payment framing: subjects are more likely to provide voluntary effort when a fixed payment of cash is framed as a gift rather than a payment for effort.

Hossain and Li (2013) run an experiment with subjects performing data entry, and in addition to varying payment levels across no or low incentives, they vary the framing of the task as either a favor to the researchers or as work. They analyze both participation rates in the voluntary work task and output levels. Under the social framing, low monetary rewards reduce participation compared to no rewards, but no such effect is found under the work framing. However, they find that among subjects who participate, low incentives increase output under both framings, in contrast to Gneezy and Rustichini's (2000a) result. In other words, they find crowding out effects on participation levels, but not on effort conditional on participation.

2.1.2. Negative evidence on crowding-out

Some recent lab experiments have not replicated the result that low piece rates reduce effort compared to no piece rate. Takahashi et al. (2016) find higher effort under low piece rates than no payment in both an interesting and a mundane task. Pokorny (2008) is also unable to

replicate the main result of Gneezy and Rustichini (2000a), even using an IQ test as the real effort task like they did. Similarly, Dessi and Rustichini (2015) use the Raven's Progressive Matrices IQ test and vary the levels of fixed payments and piece rates, but across all treatments subjects answer approximately the same number of questions correctly. They speculate that this form of test might depend more on ability than conscious effort. Alternatively, the intrinsic motivation might be so strong that it is difficult to crowd out; for example, subjects might be interested in feedback about their performance for self-image reasons. Many of these studies offer quite different levels of show-up payment and use significantly different values as "low" payments. Some authors appeal to reference-dependent preferences to explain how fixed payments could alter the effects of low piece rates; for example, subjects might work harder after a low show-up payment to reach an acceptable level of total earnings.

Two studies focus on comparing effort levels between piece rates earned for charity and piece rates earned for oneself, but incidentally fail to find crowding out when comparing low piece rates for oneself to no piece rate. Charness et al. (2016) offer subjects a choice to complete additional data entry after an hour-long experiment, with two treatments of interest being no piece rate and a two-cent piece rate earned for oneself. They find no significant difference between participation rates or output levels, and if anything find suggestive evidence of an increase in output with the very low piece rate. The authors suggest that subjects were intrinsically motivated despite the boring task because subjects thought they were helping the experimenters analyze data from another study, making the lack of crowding out at very low piece rates more perplexing. Similarly, Yang et al. (2014) use a real-effort task of finding pennies in a group of coins, and the option to keep the pennies for oneself (or donate them to charity) results in higher output than a treatment with no piece rate.

Real-effort experiments conducted in online labor markets such as Amazon's MTurk have generally failed to find negative effects from low piece rates, although most of the tasks involved would not be considered intrinsically motivating; see Buhrmester et al. (2018) for a recent literature review on the use of MTurk in experimental economics. Horton et al. (2011) offer subjects \$0.01, \$0.05, \$0.15, or \$0.25 to transcribe an additional paragraph and analyze subjects' acceptance decisions. They find monotonically increasing acceptances rates as piece rates increase, although they do not include a treatment with no incentives. Mason and Watts

(2009) recruit subjects to sort traffic pictures and complete word puzzles at varying piece rates, the latter of which could feasibly be intrinsically motivating. However, they find increasing output as piece rates increase, even for a comparison between no payment and a one-cent piece rate. This holds over all tasks and all difficulty levels of the tasks.

DellaVigna and Pope (2018) used MTurk subjects to test whether economics researchers are better than laymen at predicting the effects of various incentive levels on effort. The real-effort task was simply alternately pressing two keys on the keyboard – a task devoid of the potential for intrinsic motivation. Their treatment with a very low piece rate resulted in higher output than no piece rate; experts failed to anticipate this outcome, perhaps because they recalled the famous outcome in Gneezy and Rustichini (2000a). One fact that contextualizes these results from online labor markets is that these workers are often from very poor countries and are solely logging on to earn money. This means they could have a different perception of "low" incentives, and might have little intrinsic motivation to help researchers, in contrast to laboratory subjects.

Recent field experiments have provided interesting contexts for analyzing the effects of small incentives. For example, Carpenter and Gong (2016) paid college students to stuff envelopes for fundraising for presidential political campaigns, and randomly assigned students to matched or mismatched campaigns based on their political preferences. They conducted two "high-powered incentives" treatments with a fixed payment of $20 and piece rates of $0.50 or $1.00, and one "low-powered incentives" treatment with only the fixed payment. They found no differences across the piece rates, but there are some interesting results regarding the interaction between piece rates and intrinsic motivation. Compared to the treatment with only fixed pay, the piece rates increased productivity by 86% for subjects with misaligned political beliefs, but only by 13% for subjects with aligned political beliefs. These results suggest that if intrinsic motivation is sufficiently strong, piece rates may not impact effort much at all, whereas if intrinsic motivation is weak, piece rates can provide large productivity boosts.

Gneezy and Rey-Biel (2014) ran a field experiment on the effects of contingent and non-contingent payments for responses to consumer surveys. They varied payments from $1 to $30, with contingent payments only provided after survey completion and noncontingent payments given whether or not the survey was returned. They found that very small contingent payments

decrease response rates compared to no pay, but the opposite occurred for noncontingent payments. The authors explain that noncontingent payments may be received as a gift and exploit the respondent's reciprocity, while very small contingent payments may feel insulting.

Some economists have used econometric analysis of labor data to gather evidence on crowding out. Frey and Goette (1999) used the Swiss Labor Force Survey to analyze the effects of payments on volunteer work, finding a statistically-significant crowding-out effect: when payments were low, the hours volunteered were substantially lower than with no monetary reward. Pouliakas (2010) looked at the British Household Panel Survey to analyze the effects of incentive pay on job satisfaction. He classified the level of incentive pay as the proportion of total salary that comes from bonuses. In a panel design with individual fixed effects, he found that very low levels of incentive pay are associated with lower job satisfaction, but the effect is reversed for high levels of incentive pay. Chen (2018) used the General Social Survey to argue that small payments increase hours worked but reduce job satisfaction. In this study, very low levels of performance-based pay are associated with worse work attitudes (such as quit intentions and job satisfaction). The threshold for this effect is about 10% of total pay in the private sector and 15% in the public sector, suggesting that what is considered a "small" incentive may depend on the context and the perceptions of the specific workers involved.

In a nod to the original argument about crowding out from Titmuss (1970), some researchers have examined the effects of providing incentives for blood donations; see Frey et al. (2010), for a literature review. Mellstrom and Johannesson (2008) find evidence of crowding out particularly among women when providing cash incentives for blood donations, but they fully reverse the effect by allowing the option to donate the payment to charity. Lacetera et al. (2012) document spillover effects in which incentivized blood donations poach donors from neighboring non-incentivized blood drives. Goette and Stutzer (2019) provide novel incentives in the form of lottery tickets and free cholesterol tests. They find that lottery tickets increase blood donations particularly among less motivated donors (the cholesterol tests had no effect). In general, Frey et al. (2010) note that incentives for prosocial behavior tend to work better in anonymous situations when the incentives do not compromise the image motivations of donors to be seen as altruistic. Other research has found similar results of crowding out pro-social effort

when offering personal piece rates in addition to piece rates paid to charity (Ariely et al. 2009; Exley 2017).

Researchers have also gathered related evidence on the effects of small incentives in education (Gneezy et al. 2011 provide a general literature review of incentives in education). Levitt et al. (2016) find that a $10 incentive for improvement on a high school standardized exam did not improve scores, but a $20 incentive did. In a field experiment in China on a large online course, Gong et al. (2019) provided incentives for scoring above a certain threshold on homework quizzes. Incentives varied from zero to about four hours of wages for a Teaching Assistant at a Chinese university. The authors found no differences between no pay and very low pay, but did detect increases in submission rates and grades for the higher incentive levels. Gneezy et al. (2017) vary monetary incentives on an international standardized exam from roughly $30 to $90 and report that American students improve scores from incentives, but Chinese students did not, suggesting that the effects of incentives are sensitive even to cultural differences.

All in all, while recent research has provided considerable evidence that low performance pay can lead to poorer performance than no performance pay, there are a number of studies that do not confirm this effect, so overall the results are somewhat mixed. Whether one should expect negative effects on effort and output from offering low piece rates depends heavily on the context and even the workers themselves. Crowding out is most likely to occur in situations with strong intrinsic motivation, with very low and potentially insulting incentive pay, and when offering payment changes the framing from a reciprocal social situation to a sterile work relationship. However, even these criteria do not universally hold in the recent literature. More research is needed to delineate the relevant determinants of crowding out.

2.2. Paying too much

The intuition that paying higher wages should result in higher effort is also violated at very high levels of pay since workers can choke under the pressure to perform. This phenomenon is common in sports, where commentators discuss athletes who fail to perform in high-stakes situations. Similarly, when a worker's pay depends on output, and very large amounts of money are on the line, this pressure can cause a high level of arousal that reduces performance. However, this may be less of a risk after allowing for worker selection into high-stakes jobs:

workers who perform well under pressure are likely to end up in precisely those jobs. For example, many CEOs have millions of dollars at stake depending on their output, yet they are still able to perform at a high level.

One of the early studies in the psychology literature to explore the choking phenomenon was written by Baumeister (1984), who had subjects perform a physical task for payment and who varied the level of pressure facing the subject. He hypothesized that focusing conscious attention on a trained task can reduce performance compared to allowing automatic, unconscious systems to take over. The task involved using physical dexterity and mental focus to move two rods apart carefully so that a ball falls into a hole on a platform beneath the rods. Baumeister found that adding cash incentives, competition with other subjects, and spectators in an audience all reduced performance.

The most widely-cited paper on this topic in economics is titled "Large Stakes and Big Mistakes," by Ariely et al. (2009). The authors ran a field experiment in rural villages in India to allow them to pay very high local wages without breaking the budget. In the treatment with the largest incentives, completing the task successfully could award a subject pay equivalent to multiple months of typical income. The authors initially hypothesized that different types of tasks would be more vulnerable to choking under large incentives than others, so they included various tasks involving memory, motor skills, and creativity. It turned out that across all types of tasks, performance was reduced when moving from medium incentives to very large incentives.

At least two other experimental papers have found reductions in effort at high piece rates, although neither appealed to "choking" as an explanation and the piece rates were not nearly as relatively high as in Ariely et al. (2009). Around the same time, Pokorny (2008) published a paper finding similar reductions in performance under large incentives. In her laboratory experiment, subjects performed two real effort tasks: an IQ test and a mundane real-effort task involving counting the frequency of a particular number within a block of random numbers. Although the incentives offered were not as high relative to typical income for these subjects as in Ariely et al. (2009), Pokorny finds a significant reduction in performance between the low and high piece rates in both tasks. She appeals to the theory of reference dependence to explain the findings rather than considering the effects of pressure on performance. More recently, Takahashi et al. (2016) also found an inverse-U shaped relationship between pay and

performance, with a significant reduction in performance when offering high piece rates, but only in one of two tasks. They implemented an interesting puzzle-game task and an uninteresting task involving clicking circles that appeared on the screen. Only in the uninteresting task did the authors find a significant reduction in performance when moving from medium to high incentive pay; the interesting task aligned with the standard predictions of theory.

Evidence on choking under pressure has been gathered from contexts outside the typical work environment as well. For example, Beilock and DeCaro (2007) show that certain high-ability students suffer from anxiety and perform poorly on math exams when the stakes are high. Economists have examined evidence from sports, typically using econometric analysis to analyze the effect of stakes in playoff or championship games or other forms of pressure such as audience size. Cao et al. (2011) show athletes in the NBA are less likely to be successful in free throws that occur in the final few seconds of a game. Dohmen (2008) finds negative incentive effects from data on penalty kicks in soccer. Apesteguia and Palacios-Huerta (2010) also examine data from soccer and exploit the natural random variation in the order of kicks, and thus variation in the stakes and pressure to score. They find that teams who kick second are less likely to win due to the increased pressure, although Kocher et al. (2012) later extend that analysis to include more data and find only a negligible and insignificant effect.

In practice, choking under high incentive pay in the labor market may be less of a concern than some of these results imply because of self-selection of workers into certain jobs. Specifically, workers who perform well under high stakes are more likely to apply and thrive in exactly those kinds of positions. Macera (2017) runs a lab experiment showing that with the addition of experience in the task and self-selection, negative productivity effects of high incentives are greatly reduced. Using a math-related real-effort task, she initially identifies the fraction of subjects who perform worse under high incentives than medium incentives. She then shows that this fraction is reduced by two treatments: one that allows for significant practice in the real-effort task before performing the task for pay, and another that provides a detailed description of the task to subjects before they decide whether to enroll in the study. Both treatments reduce the fraction of subjects who perform worse under high incentives by about half relative to the baseline. Because practice and selection are common features of the real labor market, this should reduce concerns about large incentives for many employers.

Evidence on selection has also been gathered in applied data analysis. Coates et al. (2009) examine the second-to-fourth digit ratios, calculated by comparing the length of the second and fourth fingers on one's hand, of high-frequency financial traders. Although subsequent research has called digit-ratio conclusions into question, they hypothesize that because the finger ratio is correlated with testosterone levels, it could predict competitiveness and the ability to perform well under pressure. Indeed, they find that the ratio is associated with long-term profitability of these traders. Relatedly, some researchers have found evidence that women are more susceptible to this effect than men, which could partially explain the greater prevalence of men in jobs with high-powered incentives; see Azmat et al. (2016) for evidence from high stakes testing among high school students in Spain and Paserman (2007) and Cohen-Zada et al. (2017) for evidence from tennis matches.

The evidence on negative incentive effects from laboratory experiments is somewhat sparse because offering extremely high levels of incentive pay to subjects is prohibitively expensive. Ariely et al. (2009) solved this problem by recruiting subjects in poor rural villages, and others have gathered evidence from various contexts such as sports and education. However, in real labor-market situations, concerns about workers choking under high incentive pay are lessened by practice with job tasks and self-selection into high-pressure jobs. Still, it remains a consideration in certain labor contexts and an important example of the potentially non-monotonic effects of incentive pay.

3. Social-preference motivations

The workplace naturally exposes an individual worker to a rich social environment with a constellation of norms, organizational procedures, and interactions with co-workers. This section considers applications of social preferences. For example, a worker may want his boss to treat his co-workers fairly; this worker's inequity aversion may be a preference for equal payoffs among his group of co-workers. A worker who is paid a high wage may reciprocate by providing his boss with extra effort; such reciprocity is often referred to as a gift-exchange between the principal and agent. Many such examples of other-regarding behavior are common in the workplace.

This section begins with a theoretical discussion of social preferences, alongside some seminal experimental results. These fundamental models are then successively applied to various other-regarding motivations. Specifically, this section considers distributional and reciprocity social preferences, and competitive preferences. The predictions of the models and whether social preferences generalize to the field are also discussed. This brief survey of selected topics is far from comprehensive, as the literature on other-regarding preferences is vast. Bolton and Chen (2018) and Cooper and Kagel (2016) provide thorough treatments of this literature. Charness and Kuhn (2011) focus on a broader array of laboratory results, while List and Rasul (2011) focus on evidence from field experiments.

The topic of social preferences has received a great deal of attention in economics and in the broader social-science literature, particularly in the past 15-20 years. A large—and still growing—literature has documented the existence of social preferences, their social and economic implications, and the various conditions under which they are activated or suppressed. The essential content of social preferences is that people deliberately sacrifice money or other material resources to help or to hurt other people, to establish equity or equality, or to increase the economic surplus that is available for the whole group. The fact that people do not simply maximize their own material payoff has potentially far-reaching consequences for both theory and practical applications in the workplace, between the worker and co-workers, as well as between the worker and his employer.

3.1. Distributional models

Why would people choose to sacrifice their own money or resources? One way to formalize this is to consider that one puts some weight on the payoffs of another person or persons. This weight could reflect simple altruism or a preference for distribution amongst the people in one's reference group. Sacrifice could also reflect a response to actions taken by another person or persons. An action thought to be overly selfish and inconsiderate might generate negative feelings, while one considered to be surprisingly generous might generate positive feelings; the weight on the actor's well-being would be adjusted accordingly.

Models of social preferences generally fall into one of several categories. Consequential models presume that people care solely about the distribution of payoffs. The simplest approach is altruism, whereby one puts a particular weight on the payoffs of another person. This weight

may lead to different actions, depending on the price-effectiveness of sacrificing one's own money. Other models involve remedying differences in material payoffs, helping low-payoff people, or increasing the total payoff for the people affected by the choice.

The primary consequential models are those of Bolton (1991), Fehr and Schmidt (1999), and Bolton and Ockenfels (2000). Bolton (1991) develops a model in which people care about their own money, but don't like to have less than others. The latter two models also consider that receiving more than others may bother an individual. The heart of these models is that people may trade off money to make material payoffs closer together. The simple and tractable model from Fehr and Schmidt (1999) is particularly instructive, with functional form

$$U_i(x) = x_i - \alpha_i \frac{1}{n-1} \sum_{j \neq i} max \ |x_j - x_i, 0| - \ \alpha_i \frac{1}{n-1} \sum_{j \neq i} max \ |x_i - x_j, 0|,$$

where it is assumed that $\beta_i \leq \alpha_i$ and $0 \ \leq \beta_i < 1$. Here people are more bothered by being behind than by being ahead, whereas the model from Bolton and Ockenfels (2000) assumes that one's disutility from unequal payoffs is unaffected by whether one is ahead or behind.

Another consequential model is the distributional projection from Charness and Rabin (2002). This model reflects both Rawlsian and utilitarian preferences in that one has a preference for increasing the lowest payoff in the reference group as well as a preference for increasing the total payoff of these people. Specifically,

$$V_i(\pi_i, \pi_2, \dots, \pi_N) \equiv (1-\lambda)\pi_i + \lambda[\delta[\pi_i, \pi_2, \dots, \pi_N] \ + (1-\delta)(\pi_i + \pi_2 + \cdots + \pi_N)],$$

where $\lambda \in [0,1]$ measures how much i cares about pursuing social welfare versus own self-interest and $\delta \in [0,1]$ shows the trade-off between the Rawlsian and utilitarian components. The social-efficiency concern is an essential feature missing from other social-preference models. Figure 1 depicts these models for a third party with a fixed payoff x, as taken from Charness and Rabin (2000).

Figure 1: Two models of distributional preferences

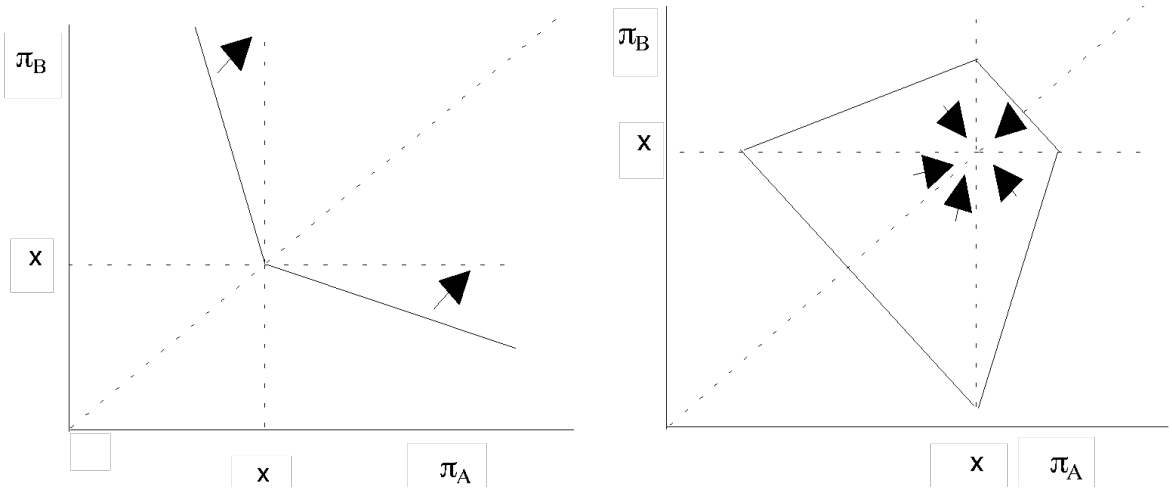

Charness and Rabin Fehr and Schmidt

These diagrams clearly show how one's preferences differ across these models. With Charness and Rabin preferences and holding one's own payment constant at x, one's utility increases as the other party's payoff increases, regardless of the direction of the difference. This illustrates the preference for efficiency that was mentioned above and is considered the hallmark of the model. By contrast, one is happiest with Fehr and Schmidt preferences when the other person's payoff is exactly the same, at x. So the models really differ with respect to how one feels about having more than someone else. This Rawlsian component of the Charness and Rabin (2002) model is an important aspect that parallels the other distributional models. In experimental tests, people were asked to choose between payoffs of (750,375,x) or (400,400,x) for (Person 1, Person 2, Self), and 54% chose the option with higher total payoffs. However, when the choice was between payoffs of (1200,0,x) or (400,400,x), only 18% chose the option with the even higher total payoffs (the value of x was disclosed only at the end and was 500). These results suggest preferences for efficiency as well as increasing the payoff of the individual with the worst outcome (Rawlsian preferences).

3.2. Competitive preferences

Of course, social preferences will not manifest in all economic environments. One is considerably less likely to make financial sacrifices in vain, as when there are powerful market forces in effect. To this effect, Fehr and Schmidt (1999) point out "It is a well-established fact

that in a broad class of market games prices converge to the competitive equilibrium [Smith, 1982; Davis and Holt 1993]." A striking example of this is the "demand game" in Roth et al. (1991). Here, nine proposers each make a simultaneous proposal to a single responder regarding the share that the proposer wishes to retain, and then the responder chooses which (if any) to accept. Proposers make very small demands and so are effectively willing to accept very small shares, unlike the results in the very similar ultimatum game. There is little point in standing on the beach and imploring the waves to stop crashing on the shore. But many economic environments feature a high degree of interpersonal interaction on a small scale, and this is where social preferences will be most in effect.

There is also a darker side to distributional social preferences. While most individuals appear to favor equalizing financial payoffs or social-efficiency considerations, other individuals prefer to come out ahead of others or perhaps simply enjoy decreasing the payoffs of others. In terms of our encompassing social-distribution model, this corresponds to the case where one's utility always decreases when the payoff of another person increases, regardless of whether the other person is ahead or behind the individual. While this is relatively rare in some circumstances, it is more common in cases where competition is highlighted. Perhaps the first evidence for competitive preferences is presented in Offerman et al. (1996), using the ring test (Liebrand 1984). The results indicate that about 15% of the subjects are willing to sacrifice money to hurt the other person rather than, for example, taking the highest own payoff and assigning a zero payoff to their counterparts.

An additional dark side to competitive preferences is seen in Charness, Masclet, and Villeval (2014). Participants could perform simple number-letter coding tasks or just chill and read magazines, receiving a fixed payment regardless of their output. Thus, there is no financial incentive to do coding tasks or to work hard. But one (in anonymous groups of three) learned after each period whether one was first, second, or last in one's group in terms of productivity. Participants could pay with their own money to either purchase additional output or to sabotage the output of others in the group. In fact, many people paid to do so, particularly when it seemed likely that this would result in an advance in one's rank in the group. This can only be explained by competitive preferences.

3.3. Intentions and reciprocity

While individuals clearly have preferences over their own payoffs and those of others, they may also care about others' intentions and motives. Whereas the distributional models of Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) only depend on the distribution of payoffs, more complex and non-consequential models incorporate beliefs or expectations. Models of intentional reciprocity take into account how selfish or generous one considers another party's choice to be and how motivated people are to reciprocate the perceived intentional actions of others.

Incorporating (intentional) reciprocity preferences is a considerably more challenging problem, but nevertheless has been addressed in several models. Rabin (1993) developed a model in which one player wishes to increase or decrease another player's payoffs based on her beliefs about whether the other player is treating her fairly.

$$U_i(a_i, b_j, c_i) \equiv \pi_i(a_i, b_j) + \tilde{f_j}(b_j, c_i) \cdot [1 + f_i(a_i, b_j)],$$

where a_i represents the strategy of player i, where b_i represents the player j's belief about what player i is choosing, and where c_i represents player i's belief about what player j believes player i's strategy is. The essential implication of this model is that people will sacrifice money to help those who they believe are being kind to them and will sacrifice money to hurt those who are being unkind to them. The kindness function $\tilde{f_j}(b_j, c_i)$ reflects beliefs about how kind one believes that the other player is being, which is defined by the location within a range of possible Pareto-efficient payoffs.

Dufwenberg and Kirchsteiger (2004) modify and extend that model to sequential games. Falk and Fischbacher (2006) also consider sequential games and combine reciprocity of the Rabin (1993) form with Fehr and Schmidt (1999) inequity aversion. Levine (1998) models reciprocity by assuming that one's preference for changing another's payoff depends on one's beliefs about the other's inherent degree of altruism. Finally, Cox et al. (2007) develop a non-equilibrium model that combines a form of distributional preferences with reciprocity considerations; one's emotional state affects one's degree of willingness to trade off own money for helping or hurting others, and in which this emotional state reflects relative payoffs and the kindness or unkindness of others.

Charness and Rabin (2002) includes negative reciprocity in two ways: First,

"misbehavior," which is determined endogenously by the group's views, diminishes the weight one places on the miscreant's payoff in the utility function. At the limit, there is no weight on this payoff in terms of either social efficiency or the minimum payoff. To begin to incorporate reciprocity, consider a strategy profile $s \equiv \{s_1, s_2, s_3, \ldots, s_n\}$ as well as a demerit profile $d \equiv \{d_1, d_2, d_3, \ldots, d_n\}$, where $d_k \in [0,1]$ for all k. The demerit profile d can be interpreted roughly as a measure of how much Player k deserves, where the higher the value of d_k, the less others think that Player k deserves. With this interpretation, we define players' preferences as a function of both their underlying "social-welfare" preferences and how they feel about other players:

$$U_1(s,d) \equiv (1-\lambda)\cdot\pi_i + \lambda\cdot[\delta min\,[\pi_i, min_{m\neq i}\{\pi_m + bd_m\}]+$$
$$(1-\delta)\cdot(\,\pi_i + \textstyle\sum_{m=i} max\,[1-kd_m, 0]\pi_m) - f\sum_{m=i} max\;d_m\cdot\pi_m)]$$

where b, d, and f are non-negative parameters of the model. The key new aspect to these preferences is that the greater is d_j for $j \neq i$, the less weight that Player i places on Player j's material payoff. Hence, these preferences say that the more that Player i feels that Player j is being a jerk, the less that Player i wants to help him. Second, when the parameter f is positive, Player i may in fact wish to hurt Player j when Player j is being a jerk.

Finally, the notion of guilt aversion takes into account that one may feel guilt to the extent that one disappoints the expectations of others (simple guilt) or believes that the other party (guilt-from-blame) will blame her. Charness and Dufwenberg (2006) show that social-preference-dependent behavior is strongly affected by non-binding free-form communication ("cheap talk") between anonymous parties, although according to the main social-preference models (whether consequentialist or reciprocity-based), this form of communication should have no effect in their set-up.

Figure 2: The extensive form of the hidden-action game

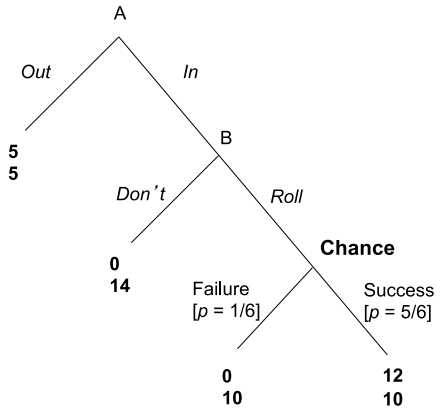

As depicted in Figure 2, the principal ("A") chooses whether to enter into a contractual relationship with the agent ("B"). If not, then each person would receive $5. If so, then B chooses whether to exert effort ("Roll") or not ("Don't"). Since effort costs $4, a B with standard (purely-self interested) preferences should choose not to exert effort. A should realize this and if s/he expects B to be selfish, A should choose "Out", leading to a poor social outcome relative to that expected from (In, Roll). In the communication treatment, B can send a free-form message to A.

Social preferences suggest that some B's will choose Roll and that some A's will choose In, for reasons that could include social efficiency (for either A or B), or minimizing the difference in payoffs or maximizing the minimum payoff (for B). Indeed, even without messages the (In, Roll) outcome occurs 20% of the time. However, communication increases this rate to 50%. Furthermore, if we consider only messages containing statements of intent ("promises"), this rate becomes 67%. Both A and B play more cooperatively when a promise is made, and the notion of guilt aversion is introduced as a driving force. The more that B's believes that A is expecting B to choose Roll, the more that B feels guilt by choosing Don't, and guilt carries some disutility (the degree of sensitivity to guilt varies across individuals). In fact, it is found that that promises affect beliefs and that there is a strong correlation between B's beliefs about A's expectations and the likelihood that B chooses to play cooperatively. This belief-dependence of social preferences thus opens up new avenues for thinking about the design of communication pathways in organizations and society more broadly. Guilt aversion may be

relevant for understanding strategic interaction in a variety of settings, and that it may shed light on the role of language and social norms in these contexts.

Note that both intentional reciprocity and guilt aversion incorporate higher-order beliefs.

3.4. Model predictions and evidence

It is natural to wonder about the extent to which the social-preference models mentioned explain the experimental data and, indeed, behavior in the field. There are many cases in which all or nearly all of these can explain the data, but often for different reasons. For example, offers of 30% or less in the classic ultimatum game are quite often rejected, meaning that responders are willing to sacrifice money to reduce the proposer's payoff. According to the Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) models, the motivation for doing so is to eliminate the disparity in the payoffs of the two parties. In all of the models containing kindness-based choices, as well as the Charness and Rabin (2002) model, rejections are instead based on negative reciprocity, punishing the deliberate choice of an offer that the responder feels is unfair.

Brandts and Solà (2001), Charness and Rabin (2002, 2005), and Falk et al. (2003) report results from mini-ultimatum and mini-dictator games and find that the rate at which people reject a low offer depends on the alternative choices that could have been made or the path by which one has arrived at the choice. In the latter study, people reject a 20% offer only 8.9% of the time when the only available alternative proposal was to offer 0%. Charness and Rabin (2002, 2005) find that 0 of 67 people choose a mini-dictator option giving each person $0 instead of $2 for the dictator and $8 for the recipient, an unusual 100% experimental result. A major weakness of consequential models such as Fehr and Schmidt (1999) is that choices depend solely on outcomes; this is clearly falsified by the examples above. And yet including reciprocity considerations may not always be necessary and these definitely come with the cost of some loss of tractability.

Consider the following three examples where the models make different predictions. In one decision task in Charness and Rabin (2002), a dictator chooses whether another person receives $4 or $7.50 and will receive $4 regardless of her choice. According to the Fehr and Schmidt (1999) model, she should choose $4. The efficiency element of the Charness and Rabin (2002) model indicates that one should choose $7.50. In fact, 69% choose $7.50. Charness and Grosskopf (2001) have a decision task in which a dictator who receives a fixed sum of 600 units

chooses whether another person receives 900 or 600 units; 66% choose 900. In a second decision task in that paper, an individual who receives a fixed sum of 600 can choose any amount between 300 and 1200, inclusive, for the other person. Seventy-four percent chose 1200, 10% chose 600, and 9% chose less than 600 (perhaps the most amusing choice was 599). Note that choosing less than 600 is only consistent with competitive or spiteful preferences.

Examples of sacrifice for punishment purposes are abundant in the literature. These include rejections in the ultimatum game and punishment in the public-goods game (Fehr and Gächter 2000, 2002), offering classic examples of negative reciprocity. Charness (2004) finds that effort choices in response to a deliberate choice of a low wage are significantly lower than when this same low wage is chosen at random (also see Blount 1995). However, positive reciprocity, in the sense of responding more favorably when another party has deliberately made a beneficial choice in one's favor, is difficult to observe in its pure form, disentangled from other mechanisms. Note that simply providing costly non-minimal effort or returning money to the sender in the Berg et al. (1995) investment game is not necessarily positive reciprocity, as these can be readily explained by preferences over distribution or a preference for efficiency. Some clean support for positive reciprocity can be found in Charness and Levine (2007), where one can arrive at a particular wage by either the choice of a high wage and bad luck or the choice of a low wage and bad luck. Gneezy and List (2006) also offers evidence of positive reciprocity. When the wage is changed to be higher than advertised and expected, workers produce more. However, this effect is evanescent. Yet the surprise is not that the effect goes away, but is instead that there is any positive reciprocity at all.

While reciprocity is difficult to isolate in the lab, studies in the field are even further complicated with confounding factors such as reputation and other repeated-game elements. What looks like positive reciprocity may in fact be favor-trading; negative reciprocity may be strategic for future interactions. One advantage of laboratory experiments is that they can cleanly isolate the pure effect. Still, there is clear evidence of negative reciprocity in the field as well. An unhappy worker might commit sabotage at the workplace, willingly taking on costs that might include the loss of his job. But observing clean positive reciprocity in the field is difficult and thus rare.

Overall, the consequentalist models do not do as good a job of organizing the

experimental data as do models that consider the path by which one arrived at the choice. Negative reciprocity is a strong feature of the environment. Furthermore, even the distributional element alone of the Charness and Rabin (2002) appendix model greatly outperforms that of the difference-aversion models. This has been established clearly established in papers such as Charness and Rabin (2002) and Engelmann and Strobel (2004).

3.5. Respect and symbolic rewards

While financial incentives are the quintessential motivator, whether self- or other-regarding in nature, workers also place value on non-financial facets of labor supply. Indeed, the ability of low wages to crowd out intrinsic motivation (as discussed in Section 1) necessarily depends on the importance of intrinsic value itself. Ellingsen and Johannesson (2007) discuss the importance of being paid respect in the workplace. Such respect can take on many forms--examples include trust, awards, recognition, and other symbolic rewards. These non-financial motivators often depend on second-order beliefs; for example, a worker places value on the beliefs that her employer believes that she is competent.

Masclet, Noussair, and Villeval (2003) consider expressions of disapproval as an alternative to monetary sanctions in the VCM. One could assign sanctioning points to another player. A sanctioned player might feel bad from this indirect punishment but would not suffer any monetary loss. In fact, both monetary and non-monetary punishments were about equally effective in the earlier periods; however, monetary punishments had a substantially greater effect on contributions than did non-monetary punishments as time passed, perhaps because people became less sensitive to the negativity expressed in non-monetary punishment points.

Eriksson and Villeval (2012) study the role of respect in an experimental labor market with repeated interaction. In this study, an employer may send a worker one to five "thumbs up," a purely symbolic gesture that is costly to the employer. The authors find that these symbolic rewards are primarily used as a coordination device to establish long-term labor relationships. Once a relationship is established, symbolic rewards are used far less frequently. Further, employers use these rewards more when the labor relationship is especially valuable (when the market has excess supply, compared to excess demand). This study also finds that respect increases stated effort provision when the labor market is balanced.

Bradler and Neckermann (2019) further investigate the value of symbolic rewards, using combinations of monetary and non-monetary gifts in two field experiments. In four separate treatments, subjects receive only a thank-you card, only money, a thank-you card and money, or a thank-you card and money with a personal touch. An increase in real-effort exertion is observed in either card-only or money-only treatments. But when a subject receives both a thank-you card and money, productivity is no different from the control without any gifts. Yet when a thank-you card is combined with money that has been personally folded into a unique shape, the productivity-enhancing effect is even stronger than in the card-only and money-only treatments. These results suggest that either monetary or non-monetary gifts can be effective at increasing productivity or perhaps reciprocity more generally. When combined, perhaps each motivator cheapens the other, resulting in no effect; yet when also combined with a personal touch, the effect of each motivator is restored, resulting in a greater effect than either has alone. Other explanations are naturally consistent with these results, and surely cultural norms are crucial to reciprocity in the field. So while most studies do not find positive reciprocity, some do, and many studies find complex interactions between motivators and the elicited responses. Bowles and Polania-Reyes (2012) provide a survey of this interaction between economic incentives and social preferences.

3.6. Do laboratory social preferences generalize to the field?

Whether the social preferences measured in laboratory experiments generalize to field settings is a question that has generated considerable controversy. Camerer (2015) provides a very extensive discussion of the issue of generalizability from the lab to the field. Notably, he argues that experimental features that could impede generalizability from the laboratory to the field (such as issues of the character of the subject pool or low stakes) are not essential for all lab experiments. He also puts forward the view that there is little evidence that typical lab features undermine generalizability. He asserts that the List (2006) data do not support the claims made in this regard and finds only small behavioral differences across the setting in List (2009). An important contribution of this paper is reporting the results of many studies designed to specifically test lab-field generalizability. These recent studies typically show a positive correlation between behavior in the laboratory and in field environments. A number of these are discussed in greater detail below.

Carpenter and Seki (2006) observe that conditional cooperation in a public-goods game predicts group fishing productivity in Japan, and Carpenter and Myers (2010) find that dictator-game allocations by Vermont residents predicts their willingness to volunteer to fight fires. Fehr and Leibbrandt (2008) show that public-goods-game contributions and patience predict limits on common-pool resource extraction by Brazilian fishermen, and Barr and Zeitlin (2010) find that dictator-game allocations made by Ugandan teachers correlates with their actual (discretionary) teaching time.

One concern is that student participants are not representative of parties making choices in natural environments. And yet studies such as Stoop et al. (2012) indicate that real-world actors also exhibit social preferences, depending on details of the environment. Recreational fishermen do cooperate in a social-dilemma game in the laboratory, but refuse to reduce their intensity of fishing in a very similar game in the field that involves actual fishing, an activity that they presumably enjoy. They conclude: "the data from these treatments suggest that the key difference between the laboratory and our field setting is the decision variable, the activity that must be undertaken in order to cooperate." So there may very well not be differences across the student and non-student populations, as much as critical differences in the richness or the per se utility from an activity.

Summing up: How much people care about social preferences in the field is not totally clear, although there is serious evidence that they do.

4. Worker responses

Standard models of labor markets usually rely on the assumption of complete contracts. However, most employment relationships are characterized by incomplete contracts, where effort is not contractible. Thus, there is much scope for social preferences to shape effort choices. This section illustrates how the existence of reciprocity and distributional preferences determine effort responses of workers to efficiency wages, wage alterations, or perceived wage inequalities. Further, evidence shows that effort responses affect the wage setting considerations of firms, in addition to having further implications for the labor market, such as an effect on unemployment. Social preferences can also enhance the enforcement of contracts. Finally, some studies show how social preferences may foster business relationships.

4.1. Gift exchange in the labor market

Akerlof (1982) proposed a model of gift-exchange to explain workers' exertion of high effort, the prevalence of wages above the market-clearing rate, and labor market unemployment. Gift exchange occurs when an employer pays its workers more than a market-clearing wage, and workers respond with more work effort than is contracted. At its essence, labor-market gift-exchange is grounded in norms regarding the actual effort workers expend, coupled with a notion of fairness within a group of workers. The difference between contracted effort and the higher fair effort is what differentiates the gift exchange from classical pure market exchange. This gap is necessary to generate reciprocity between the worker and the firm.

Incomplete contracts are often used to explain the existence of gift-exchange with efficiency wages. For example, knowledge workers' productivity may not be readily observable and thus not contractible. To ensure that these workers exert high levels of effort, firms pay these workers above a market-clearing wage. Yet in the original motivating example from Akerlof (1982), worker effort is readily observable and thus contractible, suggesting that inability to form a complete contract is not necessary for equilibrium wages above the market-clearing rate.

Fehr et al. (1993) conducted the seminal empirical test of gift-exchange, using an experiment stylized as a competitive labor market. Subjects took the role of either an employer or a worker, and employers offered wage contracts to a surplus of potential workers. If a worker accepted a wage contract, he then chose a level of stated effort for his employer. The marginal cost of effort for the worker was increasing, so that a purely self-interested worker would choose the minimal effort. The employer's payoff was multiplied by the worker's level of effort, so she can benefit greatly from a relatively inexpensive amount of worker effort. Note that a firm that believes workers are purely self-regarding will pay no more than the minimum needed. In the experiment, employers indeed provided wages in excess of worker's opportunity cost, and workers reciprocated by providing effort significantly above the minimum. The average effort was four times as high as predicted by standard theory, with minimum effort chosen in only 16% of the cases. In this study, the laboratory setting importantly permitted subject anonymity to eliminate reputation as an explanation. After the experiment the participants filled out a questionnaire, which showed that the offered wage was the major determinant for choosing a

particular effort. This indicates that workers have fairness considerations and reciprocate generous wages with higher effort choices.

These results have been widely confirmed and replicated by myriad similar studies; see for example Brown et al. (2004), Charness (2004), and Gächter and Falk (2002). Additionally, Charness (2004) investigated whether reciprocity or other forms of social preferences are responsible for the positive wage-effort relationship. Therefore, treatments varied whether wages were chosen by the employer or by an external process. He concludes that both distributional concerns and reciprocity play an important role. Overall, workers in experiments quite often respond to reasonable wages with costly effort even when this costly effort has no direct benefit to them.

4.2. Worker responses to wage raises

A purely self-regarding worker always chooses the lowest possible effort regardless of the wage offered, because providing effort is costly. Yet Akerlof and Yellen (1990) argue in their fair wage-effort hypothesis that workers respond to wage increases by providing more effort. This is because the worker has a reference wage in mind, which is shaped by various factors, such as prior wages, outside options, co-workers' wages, etc. The positive relationship between effort level and paid wage exists up to this fair wage, because a higher wage then is perceived to be fairer. Pay raises above this limit have no impact on exerted effort because one already feels fairly treated.

However, laboratory experiments are usually characterized by full information about costs or returns and provide the possibility to control for unobservable determinants of effort responses to gifts. Further, lab experiments typically do not involve the exertion of actual effort in general but simply consist of monetary transfers. As one of the first contributions to evidence in the field, Gneezy and List (2006) did two gift-exchange-type experiments with real-effort tasks. People were told that this one-time task would yield $12 per hour and last for six hours; however, subjects in the non-control treatments were informed before the experiment that they would be paid $20 per hour instead of $12. In the first half of the sessions, there is evidence of increased effort; however, the effect diminishes within a few hours and afterwards outcomes in both treatments no longer differed.

Cohn et al. (2014) find somewhat different results with a similar higher-than-advertised wage (27 CHF instead of 22 CHF). The authors find that this surprise wage increase has on average a positive and significant average effect of about 4% on productivity. It seems that a worker exerts more effort after a wage increase if and only if the wage alteration increases the perceived fairness. However, it is difficult to determine this, since participants usually do not have full information. The authors instead assessed the participants' fairness perceptions with a follow-up survey. Only workers who felt treated unfairly at the base wage showed a significantly positive effort response to a wage increase in the field. Workers who perceived their prior wage as fair did not react to the wage increase, which shows the importance of the initial wage level. A similar argument can be made regarding Gneezy and List (2006), where the baseline wage was $12 and the competitive market wage was between $8 and $10. Bellemare and Shearer (2009) also find that repeated interactions can amplify this effect. In their experiment, which took place in a tree-planting firm, workers were given a one-time wage increase. The positive effect of the wage increase is found to be stronger for workers who returned to the firm during the next planting period. Note that both these studies also provide evidence of positive wage-effort relationships in one-shot field encounters.

So there is some evidence to support the fair wage-effort hypothesis in the field. However, others such as Kube et al. (2012) found evidence suggesting that higher wages do not necessarily increase productivity and that the nature of a gift is crucial for the strength of reciprocal behavior. In their one-time experiment a monetary gift does not have a significant effect on workers' productivity; however, a monetary gift with a personal touch results in significantly higher productivity.

4.3. Worker responses to wage cuts

Akerlof and Yellen (1990) state that workers have a fair reference wage in mind. Falling below this threshold or cutting wages can decrease performance because workers may interpret it as unfair. Bewley (1997, 1999) provides insights from interviews with managers on the reasons for avoiding pay cuts. As the main determinant he emphasizes especially the possibility of damaging work morale. According to Bewley, fairness in various respects is important to good morale. One aspect of fairness is definitely payment. Workers expect firms to compensate them,

among other determinants, for qualification, tenure, performance or experience. Further, employees anticipate receiving regular pay increases, which reward good work and loyalty. Therefore, a pay cut may be interpreted as an affront or a breach of implicit reciprocity.

Experimental evidence concerning pay cuts and workers' responses to them is rather scarce. Fehr and Falk (1999) conducted a competitive double-auction experiment to examine the impact of low-wage offers on effort. Here both firms and workers can make wage offers and bids at first stage. The main treatment has incomplete contracts, where workers can choose effort after the contract has been concluded; this allows workers to reciprocate low wages with low-effort choices. Indeed, the results show that workers choose low (high) effort levels in response to low (high) wages. Workers reciprocate low wages negatively and respond with low effort levels.

In recent years researchers such as Cohn et al. (2014) and Kube et al. (2013) examined the role of pay cuts in the field experimentally, showing that indeed there is a negative effect. These field experiments complement the Lee and Rupp (2007) and Mas (2006) field studies using non-experimental data. Kube et al. (2013) hired workers to catalog books in a small library. The baseline treatment paid exactly the announced wage of €15 per hour. In the main treatment, workers were informed right before the experiment, that they would receive an hourly wage of only €10. If workers are exclusively motivated by their material self-interest, one should not observe any performance differences between the different treatments. However, if one assumes that individuals exhibit social preferences, it is expected that pay cuts decrease productivity. In fact, this wage cut decreases the average output by more than 20%. So while the effect of increased wages is insignificant or modest, pay cuts have a severe negative impact on output quantity.

To summarize, the reviewed evidence shows that there is a negative correlation between wage cuts and performance. This finding complements the evidence of increasing efforts in response to positive wage alterations as stated in the fair wage-effort theory. Considering both negative and positive responses, the question arises: which effect is stronger? Fehr et al. (2009), by comparing wage elasticities of workers' outputs in different experiments, show that, in general, cutting wages triggers stronger responses than increasing them. This means that negative reciprocity is stronger than positive reciprocity, a finding reflected in experimental papers including Charness (2004), Offerman (2002), and Charness and Rabin (2002).

4.4. Worker responses to wage inequality

The previous subsections show that workers' effort decisions depend on the perceived fairness of their wages. Survey papers such as Bewley (2002) and Card et al. (2012) have indicated that the comparison of one's wage to those of others is also a major determinant. Cohn et al. (2014) conducted a field experiment in which there were workers who performed the same tasks. In the first phase, all workers were paid the same hourly wage; in the second phase, workers in the baseline treatment were paid the same wage as before, whereas in the general-wage-cut treatment both wages were cut by 25%, and in the unilateral-wage-cut treatment only one worker's wage was cut. Workers reduced their performance in the general-wage-cut treatment on average by 15% relative to the baseline treatment and by a full 34% in the unilateral-wage-cut treatment. The wage cut in the unilateral treatment is perceived as much more unfair than in the general treatment. These results are consistent with both distributional concerns and intention-based reciprocity. Gächter and Thöni (2010) find that social comparisons matter in the domain of disadvantageous wage inequality. Clark et al. (2010) analyze how a worker's wage history, wage rank within a reference group, and average reference wage affect effort provision. The authors analyze both survey and experimental data, which suggest that workers use all of this information to determine (or gradually learn) the generosity of a given wage offer. These results suggest a nuanced interaction between firm-worker fairness, wage inequality, and worker status given by the wage rank.

Charness and Kuhn (2007) study whether co-workers' wages influence a worker's effort in the case of unequal productivities. Every firm employed two workers, a low-productive and a high-productive one; these generated unequal returns for the firm given the same effort. It was common information that workers had different productivity schedules; however, the exact details were not divulged, replicating an important feature of real-world labor markets. In one treatment, a worker knew only about her own wage, whereas in the second each worker had information about the other wage chosen by the firm. As seen before, a positive wage-effort relation is expected and shown. Furthermore, the authors argue that workers may care about co-workers' payments and exhibit inequity aversion, jealousy or envy.

As usual, effort choices depend positively on one's own wages. However, there is no significant or consistent pattern emerging for the effect of co-workers' wages. The fact that workers do not know exactly how they differ in productivity may hinder social comparison and explain this second result, because the wage inequality is not necessarily perceived to be unfair. Moreover, the authors point out that it is important to distinguish between saying one cares about others' wages, as with survey comments, and being willing to act differently because of them. A worker does not inevitably withdraw effort in this case. Still, firms tend to compress wages when both wages will be divulged.

Abeler et al. (2010) suggests that employers must account for differences in workers' performance with different wages, since high-performing workers could otherwise feel treated unfairly. The structure of their game is as follows: each principal is matched with two agents, who exert effort in stage one. Principals observe the efforts and pay a wage to the agents in stage two. In the equal-wage treatment, principals are obliged to pay the same wage to both agents. In the individual-wage treatment these wages could differ. Neither effort nor wages are contractible in either treatment. In contrast to Charness and Kuhn (2007), everyone is informed about the efforts, the wages and the resulting payoffs of all players. Abeler et al. find that workers who are paid equally provided significantly lower effort in subsequent periods than those who get individual wages, with effort levels in the individual treatment almost twice as high. It seems that agents perceive equal wages for unequal performance as unfair and thus reduce their performance afterwards. In contrast, when workers are paid individual wages, hard-working individuals continually exert high effort, and low performers change their behavior and increase their effort levels.

Both of the above experiments indicate that unequal wages might influence performance negatively only in the case of equal productivity of workers. Therefore, in some occupations, where individuals differ only slightly in productivity due to technological reasons for example, it might be optimal for employers to pay a flat wage. In the case of unequally productive workers, wages that account for these differences are necessary to avoid dissatisfied workers.

And yet equal-ability workers can be satisfied with unequal pay under some circumstances. Georg et al. (2010) conduct an extremely clever implementation of the single-principal, multi-agent model introduced by Winter (2004). Each agent works on an individual

task within a joint project, where the probability of the project's success is determined by the efforts of the agents. Critically, the technology has increasing returns to scale. The payoffs are set so that it is a dominant strategy for one agent to exert effort, it is a dominant strategy for a second agent to exert effort if the first agent exerts effort, etc. Despite the fact that equal-ability workers receive different pay for equal effort, the results show that a higher proportion of groups of workers exert full effort. In this case, individual incentives and social efficiency override the issue of inequality, consistent with the Charness and Rabin (2002) model.

5. Summary

This literature review has discussed empirical evidence on financial and non-financial motivations for worker effort. Although the neo-classical principal-agent model suggests that stronger incentives should always lead to higher effort, evidence suggests that in certain cases, this monotonic relationship does not hold. Very low levels of pay can induce lower levels of effort than no pay at all, particularly in cases where workers are intrinsically motivated to provide effort. Providing pay changes the nature of the work relationship and can crowd out intrinsic motivation from altruism or pure enjoyment of the task. Additionally, offering a low level of pay does not provide strong enough financial incentives to induce high levels of effort, thus reducing effort compared to no pay. However, this review has also discussed cases in the literature in which this crowding-out effect has not occurred, particularly in online labor markets where low levels of pay are typical. This implies that whether crowding out occurs depends heavily on the context of the work, the strength of intrinsic motivation, and what is considered a "low" level of pay for a particular population of workers.

This review has also discussed cases in which increasing pay to very high levels can reduce effort. This occurs when workers feel a high level of pressure to perform well, causing involuntary arousal that reduces performance. However, studies show that these concerns should be ameliorated in realistic labor market contexts by both self-selection into high-pressure jobs by workers who perform well under pressure, and by experience with performing work tasks under pressure.

Finally, non-financial motivations (i.e., social preferences) play a strong role in determining worker effort. In particular, social preferences for reciprocity in worker-firm

relationships often influence effort decisions. Workers often react to perceived unfairness in wages with reduced effort. In addition, perceptions of fairness depend on expected wages, changes in wages from past levels, and perhaps even on co-workers' wages. Ignoring non-financial motivations in the workplace will lead to quite sub-optimal policies; furthermore, embracing and harnessing these motivations can result in major improvements in productivity and performance.

An overriding theme across the study of these various financial and non-financial effort motivations is that the impact of each motivation depends heavily on context. Indeed, the effectiveness of an incentive policy depends on the labor market in question, the perceived fairness of the wage, and what workers consider a "low" or "high" wage, among other factors. Most past research has had a narrow focus on imposing wage policies in the lab or in the field and observing the outcomes, without taking a broad view of where and when each wage policy is effective. Further research on these topics should shed light on the contexts in which these behavioral effects are most likely to occur among different types of workers, work tasks, levels of wealth, and cultural norms regarding fairness.

References

Abeler J, Altmann S, Kube S, Wibral M (2010) Gift exchange and workers' fairness concerns: When equality is unfair. Journal of the European Economic Association 8(6): 1299–1324

Akerlof GA (1982) Labor contracts as partial gift exchange. Quarterly Journal of Economics 97(4): 543–569

Akerlof GA, Yellen JL (1990) The fair wage-effort hypothesis and unemployment. Quarterly Journal of Economics 105(2): 255–283

Apesteguia J, Palacios-Huerta I (2010) Psychological pressure in competitive environments: Evidence from a randomized natural experiment. American Economic Review 100(5): 2548–2564

Ariely D, Bracha A, Meier S (2009) Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. American Economic Review 99(1): 544–555

Ariely D, Gneezy U, Loewenstein G, Mazar N (2009) Large stakes and big mistakes. The Review of Economic Studies 76(2): 451–469

Ariely D, Heyman J (2004) Effort for payment: A tale of two markets. Psychological Science 15(11): 787–793

Azmat G, Calsamiglia C, Iriberri N (2016) Gender differences in response to big stakes. Journal of the European Economic Association 14(6): 1372–1400

Barr A, Zeitlin A (2010) Dictator games in the lab and in nature: External validity tested and investigated in Ugandan primary schools. University of Oxford, Centre for the Study of African Economies working paper

Baumeister R (1984) Choking under pressure: self-consciousness and paradoxical effects of incentives on skillful performance. Journal of Personality and Social Psychology 46(3): 610–620

Beilock S, DeCaro M (2007) From poor performance to success under stress: working memory, strategy selection, and mathematical problem solving under pressure. Journal of Experimental Psychology: Learning, Memory, and Cognition 33(6): 983–998

Bellemare C, Shearer B (2009) Gift giving and worker productivity: Evidence from a firm-level experiment. Games and Economic Behavior 67(1): 233–244

Bénabou R, Tirole J (2003) Intrinsic and extrinsic motivation. Review of Economic Studies 70(3): 489–520

Bénabou R, Tirole J (2006) Incentives and prosocial behavior. American Economic Review 96(5): 1652–1678

Buhrmester M, Talaifar S, Gosling S (2018) An evaluation of Amazon's Mechanical Turk, its rapid rise, and its effective use. Perspectives on Psychological Science 13(2): 149–154

Berg J, Dickhaut J, McCabe K (1995) Trust, reciprocity, and social history. Games and Economic Behavior 10(1): 122–142

Bewley TF (1997) Why not cut pay? Cowles Foundation discussion paper 1167, URL https://cowles.yale.edu/publications/cfdp/cfdp-1167

Bewley TF (1999) Work motivation. Cowles Foundation discussion paper 1209, URL https://cowles.yale.edu/publications/cfdp/cfdp-1209

Bewley TF (2002) Fairness, reciprocity, and wage rigidity. Cowles Foundation discussion paper 1383, URL https://cowles.yale.edu/publications/cfdp/cfdp-1383

Blount S (1995) When social outcomes aren't fair: The effect of causal attributions on preferences. Organizational Behavior and Human Decision Processes 63(2): 131–144
Bolton GE (1991) A comparative model of bargaining: Theory and evidence. American Economic Review 81(5): 1096–1136
Bolton GE, Chen Y (2018) Other-regarding behavior: Fairness, reciprocity, and trust. In: Donohue K, Katok E, Leider S (eds) Handbook of Behavioral Operations. Wiley, p 199–236
Bolton GE, Ockenfels A (2000) ERC: A theory of equity, reciprocity, and competition. American Economic Review 90(1): 166–193
Bradler C, Neckermann S (2019) The magic of the personal touch: Field experimental evidence on money and appreciation as gifts. Scandinavian Journal of Economics 121(3): 1189–1221
Brandts J, Solà C (2001) Reference points and negative reciprocity in simple sequential games. Games and Economic Behavior 36(2): 138–157
Brown M, Falk A, Fehr E (2004) Relational contracts and the nature of market interactions. Econometrica 72(3): 747–780
Camerer C (2015) The promise and success of lab-field generalizability in experimental economics: A critical reply to Levitt and List. In: Fréchette GR, Schotter A (eds) Handbook of Experimental Economic Methodology, Oxford University Press, p 249–295
Cao Z, Price J, Stone D (2011) Performance under pressure in the NBA. Journal of Sports Economics 12(3): 231–252
Card D, Mas A, Moretti E, Saez E (2012) Inequality at work: The effect of peer salaries on job satisfaction. American Economic Review 102(6): 2981–3003
Carpenter J, Gong E (2016) Motivating agents: How much does the mission matter? Journal of Labor Economics 34(1): 211–236
Carpenter J, Myers CK (2010) Why volunteer? Evidence on the role of altruism, image, and incentives. Journal of Public Economics 94(11–12): 911–920
Carpenter JP, Seki E (2005) Competitive work environments and social preferences: Field experimental evidence from a Japanese fishing community. Contributions to Economic Analysis & Policy 5(2): 2
Charness G (2004) Attribution and reciprocity in an experimental labor market. Journal of Labor Economics 22(3): 665–688
Charness G, Chen Y (forthcoming) Social identity, group behavior and teams. Annual Review of Economics
Charness G and Dufwenberg M (2006) Promises and partnership. Econometrica 74(6): 1579–1601
Charness G, Grosskopf B (2001) Relative payoffs and happiness: An experimental study. Journal of Economic Behavior & Organization 45(3): 301–328
Charness G, Kuhn P (2007) Does pay inequality affect worker effort? Experimental evidence. Journal of Labor Economics 25(4): 693–723
Charness G, Kuhn P (2011) Lab labor: What can labor economists learn from the lab? In: Ashenfelter O, Card D (eds) Handbook of Labor Economics, vol 4. North Holland, Amsterdam p 229–330
Charness G, Levine DI (2007) Intention and stochastic outcomes: An experimental study. Economic Journal 117(522): 1051–1072

Charness G, Masclet D, Villeval MC (2014) The dark side of status. Management Science 60(1): 38–55
Charness G, Rabin M (2000) Social preferences: Some simple tests and a new model. University of California working paper, URL https://escholarship.org/uc/item/46j0d6hb
Charness G, Rabin M (2002) Understanding social preferences with simple tests. Quarterly Journal of Economics 117(3): 817–869
Charness G, Rabin M (2005) Expressed preferences and behavior in experimental games. Games and Economic Behavior 53(2): 151–169
Charness G, Cobo-Reyes R, Sanchez Á (2016) The effect of charitable giving on workers' performance: Experimental evidence. Journal of Economic Behavior & Organization 131(2016): 61–74
Chen C (2018) "A little is better than zero" or "Pay enough or don't pay at all"? Evidence on the size of pay-for-performance across the sectors. Public Personnel Management 47(2): 119–143
Clark A, Masclet D, Villeval MC (2010) Effort and comparison income: Experimental and survey evidence. Industrial and Labor Relations Review 63(3): 407–426
Coates J, Gurnell M, Rustichini A (2009) Second-to-fourth digit ratio predicts success among high-frequency financial traders. Proceedings of the National Academy of Sciences 106(2): 623–628
Cohen-Zada D, Krumer A, Rosenboim M, Shapir O (2017) Choking under pressure and gender: Evidence from professional tennis. Journal of Economic Psychology 61: 176–190
Cohn A, Fehr E, Herrmann B, Schneider F (2014) Social comparison and effort provision: Evidence from a field experiment. Journal of the European Economic Association 12(4): 877–898
Cooper DJ, Kagel JH (2016) Other-regarding preferences. In: Kagel JH, Roth AE (eds) The Handbook of Experimental Economics, vol 2. Princeton University Press, p 217–289
Cox JC, Friedman D, Gjerstad S (2007) A tractable model of reciprocity and fairness. Games and Economic Behavior 59(1): 17–45
Deci E, Koestner R, Ryan R (1999) A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. Psychological Bulletin 125(6): 627–688
DellaVigna S, Pope D (2018) Predicting experimental results: who knows what? Journal of Political Economy 126(6): 2410–2456
Dessi R, Rustichini A (2015) Strong intrinsic motivation. Institute of Labor Economics discussion paper 22576
Dickinson D (1999) An experimental examination of labor supply and work intensities. Journal of Labor Economics 17(4): 638–670
Dickinson D, Villeval MC (2008) Does monitoring decrease work effort? The complementarity between agency and crowding-out theories. Games and Economic Behavior 63(1): 56–76
Dohmen T (2008) Do professionals choke under pressure? Journal of Economic Behavior & Organization 65(3–4): 636–653
Dohmen T, Falk A (2011) Performance pay and multidimensional sorting: Productivity, preferences, and gender. American Economic Review 101(2): 556–590
Dufwenberg M, Kirchsteiger G (2004) A theory of sequential reciprocity. Games and Economic Behavior 47(2):268–298

Exley C (2017) Incentives for prosocial behavior: The role of reputations. Management Science 64(5): 2460–2471
Eriksson T, Villeval MC (2012) Respect and relational contracts. Journal of Economic Behavior & Organization 81(1): 286–298
Falk A, Fehr E, Fischbacher U (2003) On the nature of fair behavior. Economic Inquiry 41(1): 20–26
Falk A, Kosfeld M (2006) The hidden costs of control. The American Economic Review 96(5): 1611–1630
Fehr E, Falk A (1999) Wage rigidity in a competitive incomplete contract market. Journal of Political Economy 107(1): 106–134
Fehr E, Gächter S (2000) Cooperation and punishment in public goods experiments. American Economic Review 90(4): 980–994
Fehr E, Gächter S (2002) Altruistic punishment in humans. Nature 415(6868): 137
Fehr E, Leibbrandt A (2008) Cooperativeness and impatience in the tragedy of the commons. Institute of Labor Economics discussion paper 3625
Fehr E, Schmidt KM (1999) A theory of fairness, competition, and cooperation. Quarterly Journal of Economics 114(3): 817–868
Fehr E, Kirchsteiger G, Riedl A (1993) Does fairness prevent market clearing? An experimental investigation. Quarterly Journal of Economics 108(2): 437–459
Fehr E, Goette L, Zehnder C (2009) A behavioral account of the labor market: The role of fairness concerns. Annual Review of Economics 1(1): 355–384
Frey B (1993) Does monitoring increase work effort? The rivalry with trust and loyalty. Economic Inquiry 31(3): 663–670
Frey B (1997) Not just for the money. Edward Elgar Publishing, Cheltenham, UK
Frey B, Goette L (1999) Does pay motivate volunteers? University of Zurich, Institute for Empirical Research in Economics working paper 7
Frey B, Goette L, Stutzer A (2010) Prosocial motivation and blood donations: A survey of the empirical literature. Transfusion Medicine and Hemotherapy 37(3): 149–154
Gächter S, Falk A (2002) Reputation and reciprocity: Consequences for the labour relation. Scandinavian Journal of Economics 104(1): 1–26
Gächter S, Thöni C (2010) Social comparison and performance: Experimental evidence on the fair wage-effort hypothesis. Journal of Economic Behavior & Organization 76(3): 531–543
Goette L, Stutzer A (2019) Blood donations and incentives: evidence from a field experiment. Institute of Labor Economics discussion paper 3580
Gong J, Liu T, Tang J (2018) Incentive Design on MOOC: A Field Experiment. National University of Singapore working paper
Gneezy U (2003) The W effect of incentives. University of Chicago, Graduate School of Business working paper
Gneezy U, List JA (2006) Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments. Econometrica 74(5): 1365–1384
Gneezy U, List J, Livingston J, Qin X, Sadoff S, Xu Y (2019) Measuring success in education: The role of effort on the test itself. American Economic Review: Insights 1(3): 291–308
Gneezy U, Meier S, Rey-Biel P (2011) When and why incentives (don't) work to modify behavior. Journal of Economic Perspectives 25(4): 191–210

Gneezy U, Rey-Biel P (2014) On the relative efficiency of performance pay and noncontingent incentives. Journal of the European Economic Association 12(1): 62–72
Gneezy U, Rustichini A (2000a) Pay enough or don't pay at all. Quarterly Journal of Economics 115(3): 791–810
Gneezy U, Rustichini A (2000b) A fine is a price. Journal of Legal Studies 29(1): 1–17
Goerg SJ, Kube S, Zultan R (2010) Treating equals unequally: Incentives in teams, workers' motivation, and production technology. Journal of Labor Economics 28(4): 747–772
Horton J, Rand D, Zeckhauser R (2011) The online laboratory: Conducting experiments in a real labor market. Experimental Economics 14(3): 399–425
Hossain T, Li K (2013) Crowding out in the labor market: A prosocial setting is necessary. Management Science 60(5): 1148–1160
Jensen M, Meckling W (1976) Theory of the firm: Managerial behavior, agency costs and ownership structure. Journal of Financial Economics 3(4): 305–360
Kocher M, Lenz M, Sutter M (2012) Psychological pressure in competitive environments: New evidence from randomized natural experiments. Management Science 58(8): 1585–1591
Kube S, Maréchal MA, Puppe C (2012) The currency of reciprocity: Gift exchange in the workplace. American Economic Review 102(4): 1644–62
Kube S, Maréchal MA, Puppe C (2013) Do wage cuts damage work morale? Evidence from a natural field experiment. Journal of the European Economic Association 11(4): 853–870
Lacetera N, Macis M, Slonim R (2012) Will there be blood? Incentives and displacement effects in pro-social behavior. American Economic Journal: Economic Policy 4(1): 186–223
Lazear E (2000) Performance pay and productivity. American Economic Review 90(5): 1346–1361
Lee D, Rupp NG (2007) Retracting a gift: How does employee effort respond to wage reductions? Journal of Labor Economics 25(4): 725–761
Levine DK (1998) Modeling altruism and spitefulness in experiments. Review of Economic Dynamics 1(3): 593–622
Levitt S, List J, Neckermann S, Sadoff S (2016) The behavioralist goes to school: Leveraging behavioral economics to improve educational performance. American Economic Journal: Economic Policy 8(4): 183–219
Liebrand WB (1984) The effect of social motives, communication and group size on behaviour in an n-person multi-stage mixed-motive game. European Journal of Social Psychology 14(3): 239–264
List JA (2006) The behavioralist meets the market: Measuring social preferences and reputation effects in actual transactions. Journal of Political Economy 114(1): 1–37
List JA (2009) The economics of open air markets. National Bureau of Economic Research working paper 15420
List JA, Rasul I (2011) Field experiments in labor economics. In: Ashenfelter O, Card D (eds) Handbook of Labor Economics, vol 4. North Holland, Amsterdam, p 103–228
Macera R (2017) Do high-powered incentives work? First aid from the labor market. Universidad de Los Andes, Chile, working paper
Masclet D, Noussair C, Villeval MC (2003) Monetary and nonmonetary punishment in the voluntary contributions mechanism. American Economic Review 93(1): 366–380

Mason W, Watts D (2009) Financial incentives and the performance of crowds. In: Proceedings of the ACM SIGKDD workshop on human computation, Association for Computing Machinery, New York, p 77–85
Mas A (2006) Pay, reference points, and police performance. Quarterly Journal of Economics 121(3): 783–821
Mellström C, Johannesson M (2008) Crowding out in blood donation: was Titmuss right? Journal of the European Economic Association 6(4): 845–863
Offerman T (2002) Hurting hurts more than helping helps. European Economic Review 46(8): 1423–1437
Offerman T, Sonnemans J, Schram A (1996) Value orientations, expectations and voluntary contributions in public goods. Economic Journal 106(437): 817–845
Paserman M (2007) Gender differences in performance in competitive environments: Evidence from professional tennis players. Institute of Labor Economics discussion paper 2834
Pokorny K (2008) Pay—but do not pay too much: An experimental study on the impact of incentives. Journal of Economic Behavior & Organization 66(2): 251–264
Pouliakas K (2010) Pay enough, don't pay too much or don't pay at all? The impact of bonus intensity on job satisfaction. Kyklos 63(4): 597–626
Rabin M (1993) Incorporating fairness into game theory and economics. American Economic Review 83(5): 1281–1302
Roth AE, Prasnikar V, Okuno-Fujiwara M, Zamir S (1991) Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. American Economic Review 81(5): 1068–1095
Stoop J, Noussair CN, Van Soest D (2012) From the lab to the field: Cooperation among fishermen. Journal of Political Economy 120(6): 1027–1056
Takahashi H, Shen J, Ogawa K (2016) An experimental examination of compensation schemes and level of effort in differentiated tasks. Journal of Behavioral and Experimental Economics 61(2016): 12–19
Titmuss R (1970) The gift relationship: from human blood to social policy. Pantheon Books, New York
Winter E (2004) Incentives and discrimination. American Economic Review 94(3): 764–773
Yam K (2013) Experimental manipulations to reduce the negative effects of money. Journal of Applied Social Psychology 43(5): 1104–1109
Yang A, Hsee C, Urminsky O (2014) Eager to help yet reluctant to give: How pro-social effort and pro-social choices diverge. Social Science Research Network working paper 2733331