In Search of Informed Discretion:

An Experimental Investigation of Fairness and Trust Reciprocity

Victor S. Maas
University of Amsterdam
v.s.maas@uva.nl

Marcel van Rinsum
RSM Erasmus University
mrinsum@rsm.nl

Kristy Towry*
Emory University
kristy_towry@bus.emory.edu

August, 2009

*Corresponding author

Acknowledgements
This paper has benefitted from comments and suggestions at the 2009 AAA Annual Meeting, the 2009 AAA Management Accounting Section midyear meeting, the 2009 EAA annual meeting and the 2008 Tinbergen Institute Conference on Human Relations, Reciprocity and Incentives in the Workplace and by workshop participants at the University of Amsterdam (CREED), Emory University, Erasmus University Rotterdam, the University of Ljubljana and Tilburg University. We would also like to thank Ivo van Amelsfoort, Leslie Berger, Martine Cools, Tim Mitchell, Gary Hecht, Arthur Schram and Sergeja Slapnicar for help and useful comments.

In Search of Informed Discretion:
An Experimental Investigation of Fairness and Trust Reciprocity

ABSTRACT

This paper investigates managerial discretion in compensation decisions in a team (i.e., joint production) setting. Specifically, we investigate the conditions under which managers tasked with allocating a discretionary bonus pool are willing to incur a personal cost to obtain *ex post* non-contractible information about the individual effort levels of of team members. Using theory from behavioral economics that incorporates preferences for fairness into the manager's utility function, we predict and demonstrate experimentally that managers' willingness to incur such a cost increases as the team's aggregate performance becomes less extreme (i.e., as the team's aggregate performance becomes a more noisy measure of individual performance). Further, using theory that incorporates preferences for trust reciprocity into the manager's utility function, we predict and demonstrate experimentally that managers' willingness to incur the cost will be greater for relatively high vs. relatively low levels of aggregate performance. The study contributes to the existing literature on subjective performance evaluation by identifying how social preferences influence managers' use of discretion in evaluation processes.

Keywords: *subjective performance evaluation, discretionary bonus pool, incentive contracting, reciprocity, fairness, trust, third-party intervention.*

Data availability: *The data from this study and the set of instructions for the experimental task are available from the researchers upon request.*

I. INTRODUCTION

Accounting research on incentives and contract design has traditionally taken an agency theory approach, deriving optimal contracts based on objective performance measures (e.g., Holmstrom 1979). In more recent years, this line of research has expanded to consider the use of discretion in performance evaluation and compensation (e.g., Rajan & Reichelstein 2006; Macleod 2003; Baiman & Rajan, 1995). The intuition behind this newer line of research is that discretion can be optimal when managers have access to non-contractible information, which supplements contractible information to provide a more complete picture of performance. For example, consider a team setting, in which individuals combine their efforts in joint production work, as is common in both manufacturing (e.g., production cells) and service (e.g., client services) settings. The team's aggregate performance, while contractible, is likely a noisy measure of individual performance. Additional (non-contractible) information, such as that obtained via supervisory observation and investigation, may be useful in filtering out that noise to distinguish individual performance.

The influence of non-contractible information depends not only on its availability, but also on managers' willingness to make use of it. However, it is well established that evaluating managers often exhibit a *centrality bias*, such that differences in performance across employees are understated (e.g., Bol, 2009a, 2009b). Relatedly, recent experimental work suggests that in allocating discretionary bonus pools, managers tend to allocate the bonus pool evenly or to overly rely on readily available contractible information (Bailey et al. 2009). These results suggest that managers may fail to fully consider non-contractible information, and it is, thus, important to understand the factors that influence this tendency. In this paper, we examine the conditions under which managers are willing to incur a personal cost to obtain non-contractible information that would supplement the information conveyed by a contractible measure of a team's aggregate

performance. Using behavioral economics theory on social preferences, we argue that managers' willingness to do so will be influenced by the outcome of the contractible measure. Thus, we add to the literature examining the interplay between contractible and non-contractible performance information (e.g., Bailey et al. 2009, Bol and Smith, 2009).

We conduct our investigation using a laboratory experiment in a team production setting with a discretionary bonus pool.[1] The experimental setting is a one-period, three-person game, with two employees and one manager. The two employees engage in a team effort, such that only their combined performance can be observed costlessly. As is common in practice, a bonus pool is funded based on aggregate performance (i.e., a contractible performance measure), and the manager is given full discretion in terms of how to allocate the pool between the two employees. Before making the allocation decision, the manager has the opportunity to obtain, at a personal cost, non-contractible information about the individual effort levels of the employees. If obtained, this non-contractible information allows the managers to unravel aggregate performance in order to make more informed allocation decisions.[2]

As expected, we find that, in general, managers are willing to forego wealth in order to obtain the information. More importantly, we find that the willingness to do so depends in two ways on the outcome of the contractible performance measure (aggregate performance). First, we find support for our hypothesis, based on the managers' preferences for fairness, that their willingness to obtain costly, non-contractible information increases as aggregate performance becomes less extreme. Second, as predicted by theories of trust reciprocity, we find that managers

[1] Our study is an economic experiment in the sense that we observe individuals' economic decision making in a controlled laboratory setting with monetary incentives (Callahan et al., 2006). However, it should be noted that we measure rather than manipulate the effort levels of the employee participants in our experiment. As the combined effort of the employees is the main independent variable in our setting, our measurement of this variable strictly speaking renders our study a quasi-experiment (Shadish et al., 2002).

[2] While throughout this paper we use the terms "employee" and "manager," our theory generalizes to any hierarchical setting with subordinate and superior managers.

are more willing to obtain costly, non-contractible information if aggregate performance is relatively high than if it is relatively low. The second result is particularly notable, because it works in the opposite direction of the effect described in the recent analytic literature (Rajan and Reichelstein, 2009).

Our study follows up on recent work that uses experimental economics and behavioral game theory to increase our understanding of management accounting issues (e.g., Drake and Haka 2008; Hannan 2005; Kachelmeier and Towry 2002). The primary contribution of this paper is an increased understanding of how social preferences affect the use of non-contractible information in discretionary bonus pool allocations (Bailey et al. 2009; Ittner et al. 2003). While it is well established in the literature that such decisions may be influenced by the manager's self-interest (Prendergast and Topel 1993) and cognitive biases (Bol and Smith 2009; Moers 2005; Lipe and Salterio 2000), relatively little is known about how social preferences affect the decisions of evaluating managers. An important implication of our findings is that managers' use of discretion is likely to involve a tradeoff among fairness, trust reciprocity, and the costs of acquiring extra information.

From a more fundamental behavioral economics perspective, a unique aspect of our study is our focus on the effect of social preferences on decision making *under uncertainty*. That is, while prior research examines the decision maker's response to unfairness, we examine the decision maker's willingness to incur a cost to prevent potential, but uncertain, unfairness. While real-world decision makers commonly face such uncertainties, we know little about how they react in these situations. Our results suggest that managers are more likely to intervene as the potential for unfairness increases, but that this effect is moderated by judgments of whether their employees have earned the right to fair treatment. We furthermore find that participants' deviations from the model of "economic man" cannot be fully explained by inequity aversion (Fehr and Schmidt

1999; Bolton and Ockenfels 1993), as the managers in our study often use their discretion to *increase* rather than *decrease* differences in employee allocations. The findings instead support models of third party punishment and the notion of "strong reciprocity" (Charness and Rabin 2002; Fehr *et al*. 2002).

This paper proceeds as follows. In Section II we provide the theoretical background of our study and develop our hypotheses. Section III describes our research design. In section IV we present the results of our study. Finally, Section V provides a discussion of our results, conclusions and directions for future research.

II. THEORY AND HYPOTHESIS DEVELOPMENT

Background

While the accounting literature often focuses on complete contracting based on objective performance metrics, the use of discretion in performance evaluation and compensation is quite common in practice (Gibbs *et al*. 2004). An important reason for the use of discretion is that contracts are typically incomplete, and discretion allows managers to consider the effects of non-contractible information. Contractible performance metrics usually provide only noisy measures of performance. For example, macro economic changes can have profound effects on contractible performance metrics such as profitability. These macro economic changes are uncontrollable and often unanticipated, and managers can use their discretion to untangle the effects of such factors from performance metrics.

The use of discretion can be particularly valuable in a joint production or team setting, in which organizational participants collaborate to produce a joint outcome. Such settings often provide contractible measures of aggregate performance. However, these contractible measures provide scant information on individual contributions to joint production. Managers must rely on

other sources of information to assess individual performance. These other sources, such as direct observation and general impressions, are generally not jointly verifiable and are thus non-contractible. However, they are still relevant, and firms have devised methods of incorporating such non-contractible information into performance evaluation and compensation decisions.[3]

A common method of allowing for discretion, which we investigate in this paper, is through the use of discretionary bonus pools. Firms commonly fund bonus pools based on contractible measures of aggregate performance, and then endow managers with discretion to allocate the bonus pools using whatever relevant information is available – both contractible and non-contractible. Baiman and Rajan (1995) show that the use of such bonus pool arrangements, which allow managers discretion over the *allocation* of the bonus pool but not its *size*, can mitigate incentive problems on the side of the manager and enable the efficient use of non-contractible information. This analytic result is confirmed experimentally by Fisher *et al.* (2005).

The usefulness of discretion depends not only on the availability of non-contractible information, but also on the evaluating manager's willingness to obtain and use this information. Investigations and evaluations have costs attached to them, as managers can invest resources, time and effort in ways that are more congruent with their direct self-interests. Therefore, managers may be reluctant to obtain and use non-contractible information about employee performance. We examine the conditions under which managers will be willing to obtain non-contractible information that allows for more informed bonus allocations. As will be explained below, we expect that managers' willingness to obtain costly, non-contractible information will depend on the outcome of contractible measures.

[3] By allowing managers to consider a broader range of information, discretion can strengthen the links between effort and compensation, essentially reducing the noise inherent to systems that rely only on contractible performance metrics. Thus, discretion can reduce risk to the employee and consequently, reduce the cost of incentive contracting (Baker *et al.* 1994).

A considerable body of analytic literature similarly considers the effect of contractible performance measure outcomes on the use of non-contractible information. Rajan and Reichelstein (2009), for example, show that in a discretionary bonus pool setting, contractible and non-contractible information may be optimally considered in a lexicographic manner, with a non-contractible performance measure used only if the contractible measure's outcome is sufficiently low.[4] Similarly, Ederhof (2008) shows that non-contractible information is optimally considered only when the contractible performance measure outcome falls outside a "normal" range.[5] These results relate specifically to discretionary bonus pools, but they are logically related to the prior literature on conditional variance investigations (e.g., Lambert 1985, Young 1986), which has considered the optimality of one-sided vs. two-sided variance investigation policies.

These analytic papers examine how the outcomes of contractible measures influence the use of non-contractible measures, within the confines of optimal contracting based on traditional principal-agent models. Specifically, the models assume that the principal has utility for wealth only. In contrast, we assume that the evaluating manager (i.e., the principal) exhibits other-regarding preferences (i.e., preferences for fairness and trust reciprocity). Thus, while the analytic literature examines the interplay between contractible and non-contractible performance measures using notions of optimal contracting, we examine this interplay based on notions of managers' social preferences. The effects of social preferences may be substantial, as existing research indicates that honesty, fairness, and reciprocity impact a wide range of individuals' economic decisions (Fehr and Fischbacher 2002; see Camerer 2003 for a review), including decisions within the domain of management accounting (e.g., Matsumura and Shin 2006; Hannan 2005; Evans *et al*. 2001).

[4] Rajan and Reichelstein's (2009) results vary depending on whether they consider a single- or multiple-agent setting, and depending on the solution concept used (Nash equilibrium versus dominant strategies).
[5] While Rajan and Reichelstein (2009) assume that the two performance measures are independently conditional on the employee's action, Ederhof (2008) assumes that they are not.

Basic setting

In our study, we consider a situation in which employees engage in a joint production effort, such that only their combined performance can be observed without cost. Thus, aggregate performance (a contractible measure) is a noisy indicator of each employee's individual contribution. The basic setting is similar to that of Baiman & Rajan (1995) and Fisher *et al*. (2005), and consists of two employees and one manager. Both the employees and the manager receive a fixed endowment. The employees independently choose how much effort (operationalized as a monetary cost) to devote to joint production, and their combined effort determines aggregate performance. That is, aggregate performance is a noiseless measure with regard to *combined* effort. A bonus pool is funded formulaically on the basis of aggregate performance. Table 1 demonstrates the relation between employee effort and the total bonus pool amount, using the parameters used in this study. Note that the total bonus pool equals 150% of the contractible performance measure (i.e., 150% of the combined effort of the two employees).

--- Table 1 about here ---

The manager has full discretion over the allocation of the pool to the two employees. Different from Fisher *et al.* (2005), the manager in our setting can obtain information on individual contributions to aggregate performance only via a costly search. Operationally, we give the manager the opportunity to purchase information about the individual effort levels of the employees. The employees know about the manager's option to obtain and use this information. However, the setting does not allow communication, and so the manager cannot commit *ex ante* to obtaining the information. The employees therefore have to make an estimate of the probability that the supervisor will obtain and use this information.

Suppose an employee expects the manager to maximize his/her monetary payoff, which

means that s/he will not be willing to incur a cost to obtain this information. In this situation, the employee can reasonably expect to receive an allocation equal to half of the bonus pool.[6] As long as the return on effort is between zero and 100% (i.e. the size of the total bonus pool is somewhere between the total cost of effort and twice the total cost of effort), employees effectively find themselves in a prisoner's dilemma, where the dominant strategy is to choose zero effort. In our setting the return on effort is 50%, in that the total bonus pool is funded as 150% of the total of the two employees' effort levels. Table 1 indicates that if the manager simply allocates half of the bonus pool to each employee, then each employee's wealth maximizing effort choice is zero. Therefore, the Nash equilibrium is an investment of zero effort for both employees.

Suppose, on the other hand, that the employee expects the manager to always obtain costly, non-contractible information on individual contributions to joint production, and then to allocate the bonus pool proportional to individual effort. In this case, the employee's optimal strategy is to maximize his/her effort, as every unit of effort returns a positive reward.

Prior research in behavioral game theory (Camerer 2003) leads us to expect most employee participants to engage in a positive level of effort, given some chance of investigation. This is important, as our experiment requires that we observe variation in employees' combined effort level, which is our independent variable. Thus, we offer this prediction not as a hypothesis, but rather as an assumption that must be met in order for us to investigate our research question. The experimental setting is described in more detail in section III.

Hypotheses

Under the assumption that the manager maximizes wealth, standard economics-based

[6] There are two reasons why in this situation the employee should expect to get half of the bonus pool. First, managers may distribute the pool randomly over the two employees, in which case the long term average allocated reward will be 0.5 times the bonus pool. Alternatively, research (e.g. Bailey *et al.* 2009) suggests that in absence of information many managers will anchor on a 50-50 split. In both cases the expected value is half the bonus pool.

reasoning suggests that in a one-period world, managers will never obtain costly, non-contractible information on individual effort, because doing so would reduce their monetary payoff.[7] Psychological theories of organizational justice (Cropanzano *et al*. 2001) and recent insights from experimental economics and behavioral game theory (Camerer 2003; Fehr and Schmidt 2003), however, suggest that individuals care about fairness and trust reciprocity. By obtaining information on individual performance and by acting on this information, the manager can ensure that each employee receives a fair allocation, thus rewarding employees who act cooperatively, punishing employees who try to free-ride, and enforcing a social norm of cooperation (Fehr and Fischbacher 2004b; Fehr et al. 2002). Studies suggest that individuals derive utility from enforcing social norms (Fehr and Fischbacher 2004a; Bendor and Swistak 2001) and that cooperation is an important social norm in human societies (Gintis 2000; Fiske 1991). Deviation from this norm is perceived to be unfair and is associated with negative emotions (Reuben & Van Winden, 2008; Van Winden, 2007; De Quervain et al. 2004). Accordingly, there is evidence of a basic human tendency to reward cooperative behavior and punish uncooperative behavior, even if this behavior was directed toward some other person (Fehr and Fischbacher 2004a, 2004b).

Three studies in particular show that independent observers are willing to incur a cost to reward and punish others (Turillo *et al*. 2002; Fehr and Fischbacher 2004b; Kahneman *et al*. 1986). First, Kahneman *et al*. (1986) report that 75 percent of their experimental participants were willing to incur a small cost to reward cooperative and punish uncooperative behavior.[8] Turillo, *et al* (2002) replicated this experiment and report a similar 73 percent of their participants choosing a

[7] Reputation effects provide one limit to such behavior in a repeating setting, and the contracting literature often turns to reputation as the self-enforcement mechanism that substitutes for court verifiability if non-contractible performance information is incorporated via discretion (Baker *et al*. 2002; Prendergast & Topel, 1993).

[8] In their experiment, participants were given the choice between (1) an equal split of $12 between themselves and someone who had acted uncooperatively in an earlier experiment, and (2) an equal split of $10 between themselves and someone who had acted cooperatively in the earlier experiment. The cooperative person had chosen a $10/$10 split between themselves and someone else instead of an $18/$2 split, while the uncooperative person had chosen to take the $18 and leave the other participant with $2.

self-sacrificing split with a cooperator. In a series of follow-up experiments, they disentangled the effects of rewarding cooperation and punishing selfishness, finding that individuals were prepared to pay for both. Finally, Fehr & Fischbacher (2004b) let their participants observe two other individuals playing prisoner's dilemma games and dictator games. In both games, they found that significant numbers of the participants were willing to incur a cost to punish a player who had behaved uncooperatively.

This prior literature thus demonstrates a willingness of third parties to sacrifice wealth in order to reward (punish) others who *are known to* have behaved cooperatively (uncooperatively). We extend this literature by addressing a setting in which the third party is a step removed from the decision to reward or punish others. Specifically, the manager in our setting does not know how the individual employees have behaved. Therefore s/he does not have any reason to deviate from an equal split of the reward, which may be considered the fairest alternative in absence of information about individual effort. However, what if such information is available at a personal cost to the manager? In accordance with the prior literature, we predict that managers will be willing to incur a cost in order to gain information that will allow them to reward or punish employees in accordance with their relative effort levels, because they will consider such an allocation to be fairer than the essentially arbitrary equal split. That is, we begin with a baseline hypothesis, suggesting that managers charged with allocating a bonus pool are generally willing to obtain costly, non-contractible information on individual contributions to group production.

H1: Managers are willing to incur a cost in order to obtain information on the individual contributions to joint production.

We furthermore expect that the willingness to obtain this information will depend in two ways on the outcome of aggregate performance (i.e., the contractible performance measure). First, we predict that the managers' willingness to obtain costly, non-contractible information will

increase as aggregate performance becomes less extreme. Second, we predict that, holding extremeness constant, managers are more willing to obtain costly, non-contractible information if aggregate performance is relatively high than if it is relatively low. We describe the theory underlying these predictions below.

With regard to the first of these predictions, as aggregate performance becomes more extreme (either high or low), it becomes a less noisy indicator of individual performance. That is, exceptionally good aggregate performance can only be the result of substantial effort of all responsible employees and exceptionally bad aggregate performance can only mean that all employees have failed to excel. In our setting, the noisiness of aggregate performance as a measure of individual performance is reflected in the number of possible combinations of effort choices by the two employees that could have led to the observed outcome. This number increases as the aggregate measure of team performance moves away from the extreme (high or low) outcomes.

The point is illustrated in Table 1. Suppose first that aggregate performance equals zero. This can only mean that both employees have chosen effort levels of zero, and so in this case, the aggregate measure is also a perfect measure of individual performance. Similarly, if the aggregate measure is 20 (resulting in a total bonus pool at the maximum of 30), this can only mean that both employees have put forth their maximum effort. At all other levels, more than one combination of effort choices by the two employees is possible. The number of possible combinations is maximized for a combined effort level of 10, with a corresponding total bonus pool of 15. Note in Table 1 that there are 11 different combinations of effort that lead to a total bonus pool of 15.

Recall that we expect managers to exhibit preferences for fairness. The fairness effect implies that managers will be more willing to obtain costly, non-contractible information on individual contributions to joint production when aggregate performance is in the moderate region

than when it is in either the high or low regions. As the aggregate performance measure moves away from the extremes, it becomes more likely that an equal split of the bonus pool will punish an employee who has invested a substantial amount of effort and reward the uncooperative behavior of another employee. If, as the literature suggests, individuals want to ensure just desserts (Carpenter *et al*. 2004; Fehr and Fischbacher 2004b; Mohtashemi and Mui 2003), managers will be more willing to incur a cost for information on individual contributions to joint production as the aggregate measure becomes less informative about individual efforts. Based on this reasoning, we make the following formal prediction.

> **H2:** The cost that managers are willing to incur in order to obtain information on individual contributions to joint production increases as the aggregate performance measure moves away from either of the extremes (i.e., as the aggregate performance measure becomes a more noisy measure of individual performance).

We now turn to the second way in which we expect the willingness to obtain costly, non-contractible information to depend on the outcome of the contractible performance information (i.e., aggregate performance). We expect that managers will be more willing to obtain costly, non-contractible information on individual performance when the employees have performed relatively well in the aggregate than when they have performed relatively poorly in the aggregate. We base this prediction on the notion of trust reciprocity.

When employees choose to invest effort, knowing that their return depends on the managers' discretionary allocation, they signal trust that the manager will provide them with a fair return. The notion of trust reciprocity suggests that managers will derive utility from repaying this trust by ensuring that employees indeed earn such a fair return on their effort. This prediction is consistent with the considerable literature suggesting that individuals are willing to repay other individuals who have made trusting decisions (Hannan, 2005; Fehr and Schmidt 2003; Dufwenberg and Gneezy 2000; Berg *et al*. 1995). In contrast, if the aggregate performance

measure indicates that both employees have invested limited amounts of effort, the manager will feel less inclined to obtain costly, non-contractible information in order ensure that the bonus pool allocation fairly reflects individual contributions. In this case, neither employee has placed great trust in the manager, and so the manager is more likely to allocate half of the bonus pool to each employee or to make a random allocation.

To illustrate, again refer to Table 1. Suppose that the aggregate performance measure equals 18 (resulting in a bonus pool of 27.0). This level of aggregate performance is associated with three possible effort combinations. It is certain, however, that both employees have invested substantial effort, as each employee must have chosen an effort level of at least 8 to achieve this outcome. Now compare this situation to one with aggregate performance of 2 (resulting in a bonus pool of 3.0). Again, this aggregate performance can only have come about by three different effort combinations. Yet, in this case both employees have clearly acted uncooperatively, as neither employee could have chosen an effort level greater than 2. Due to preferences for trust reciprocity, we hypothesize that managers will be more willing to obtain costly, non-contractible information on individual contributions to joint production in the former case than in the latter.

> **H3:** The cost that managers are willing to incur in order to obtain information on individual contributions to joint production will be greater for relatively high than for relatively low levels of aggregate performance.

On the surface, this prediction may seem to conflict with recent work that suggests a strong tendency for third parties to punish others who are perceived to violate a social norm (De Quervain et al. 2004, Fehr and Fischbacher 2004b, Hannan *et al.* 2009). That is, one may interpret H3 as predicting a stronger preference for rewarding than for punishing employees. Recall, however, that the funding of the bonus pool is based on the team's aggregate performance. Thus, the team *as a whole* will be automatically punished for poor performance and rewarded for good performance. The role of the manager's discretionary allocation is to distinguish individual

performance. Thus, even when the team's aggregate performance is relatively high (low), the manager's motivation for obtaining and using costly, non-contractible information could include punishing (rewarding) the lower (higher) performing employee.

It is important to discuss the pattern of results that are predicted by H2 and H3. H2, which is based on the fairness argument, suggests that managers will be more willing to obtain costly, non-contractible information as the aggregate measure moves away from either of the extremes. H3, which is based on trust reciprocity, suggests that managers will be more willing to obtain such information as the aggregate measure increases. When combined, the two hypotheses suggest a non-linear relation between the aggregate performance measure and the manager's willingness to obtain costly information. This pattern is depicted in Figure 1, Panel A, and described in further detail below.

Using the actual parameters from the study, the aggregate measure of combined effort ranges from 0 to 20. As the aggregate measure increases from low levels of effort to moderate levels (i.e., 10), both the fairness (H2) and the trust reciprocity (H3) arguments suggest that the manager's willingness to obtain this information will increase. However, as the aggregate measure increases from moderate to higher levels, the two arguments offset each other, with the fairness argument (H2) predicting a decrease in the willingness to obtain information and the trust reciprocity argument (H3) predicting an increase. Thus, the slope of this line becomes less steep, reflecting the positive effect of trust reciprocity (H3) and the offsetting negative effect of fairness (H2). Depending on the relative sizes of these effects (which we do not predict), the slope may remain positive or become negative. Importantly, if the slope becomes negative (indicating a stronger effect for fairness than for trust reciprocity), we would expect this negative slope to be less steep than the positive slope to the left of 10. That is, the trust reciprocity effect will create an asymmetry in the pattern.[9]

[9] To understand the pattern, first assume that the managers' decisions were driven solely by fairness concerns and that trust reciprocity played no role. The fairness effect we hypothesize is driven by the fact that aggregate performance is

--- Figure 1 about here ---

III. RESEARCH DESIGN

Experimental Design

We conducted a three-person laboratory game that combines elements of a one-shot prisoner's dilemma and the trust game of Berg *et al*. (1995). Notably, our study involves no experimental manipulation. Rather, the independent variable (aggregate performance) is measured. Participants interacted anonymously with each other through a computer network in the laboratory. The game was programmed using the software package Z-tree (Fischbacher 2007). Table 2 lists all design parameters and variables, and Table 3 provides a time line of the ten stages in each round of the experiment. All monetary amounts were denoted in an experimental currency (Lira) which has a value of 0.5 Euro.[10] The employee endowment was 10 Lira and the manager endowment was 15 Lira.

--- Table 2 about here ---

--- Table 3 about here ---

Following the experimental economics tradition, we operationalize employee effort as a monetary cost (e.g., Brüggen and Strobel 2007). That is, each employee must choose an effort level in the range of 0 to 10, which results in a payment of 0 to 10 Lira from the employee's initial

a noisy measure of individual performance, and only by obtaining information on individual contributions to joint production can the manager ensure that rewards will be proportional to contributions. Because the noise in aggregate performance as a measure of individual performance is symmetric around 10 (e.g., a combined effort of 8 or 12 could both by caused by exactly 9 different combinations of effort), we would expect to see a perfectly symmetric relation between the aggregate performance measure and the managers' price offers, with the offers increasing up to an aggregate performance measure of 10 and then decreasing at the same rate beyond that point. Suppose, on the other hand that the managers' decisions were driven solely by trust reciprocity and that fairness played no role. The effect of trust reciprocity is that managers will be more willing to incur a cost to obtain information on individual contributions to joint production as the aggregate performance measure increases. If this were the only effect in play, the noise in the individual performance measure would be irrelevant, and we would expect no change in the slope at 10. Rather, we would expect a consistent and monotonic increase in price offers as the aggregate performance measure increased from 0 to 20. If both fairness and trust reciprocity are playing roles, we would expect a reduction in the slope at the pivot point of 10, but a lack of symmetry, such that the absolute value of the slope to the right of 10 will be lower than the absolute value of the slope to the left of 10.

[10] At the time the experiment was run the exchange rate was 1.55 US Dollar for 1 Euro.

endowment. This operationalization is consistent with the notion of effort as a construct that is controllable by the employee, creates negative utility, and results in an increase in expected output (Baiman 1982).

Notably, the managers' pay is not a function of aggregate performance. We make this design choice to isolate the effects of the social preferences of interest. Specifically, in our setting, if the manager has utility only for wealth (as would be assumed in the traditional agency theory model), then we would expect him/her to be unwilling to obtain costly, non-contractible information. Therefore, if we observe manager participants obtaining the non-contractible information, we can conclude that their decisions are driven by social preferences. Further, if the manager's payoff had depended on aggregate performance, the effects of fairness and reciprocity would be confounded by the fact that the cost of non-contractible information would be relatively cheaper for higher levels of aggregate performance. Importantly, this design choice thus biases against finding the hypothesized effects.

We use a Becker-DeGroot-Marschak (BDM) mechanism (Becker *et al.* 1964; Bohm *et al*. 1997) to assess the managers' willingness to obtain costly, non-contractible information about individual contributions to joint production. Specifically, participants were informed that such information was available for a price, which was unknown and would be determined randomly by the computer. After the employees had decided on their effort levels and all three players had learned the value of the available bonus pool, the managers made a price offer for the information. This price offer is the main dependent variable in our study, as it represents the manager's willingness to pay for this information. After the managers made their offer, the actual price was randomly determined. Managers only obtained the information, at the *actual* price, if the offer was at least as high as the actual price. If the offer was below the actual price, the manager did not receive the information and paid nothing (Becker *et al*. 1964).

All participants knew the (uniform) probability distribution of the actual price of the information. They were informed at the start of the session that the price of the information varied between 0 and 5 Lira, such that for an offer of 0 a manager was certain not to receive the information, while for an offer of 5 Lira s/he was 100 percent certain to receive the information. The chance of obtaining the information increased linearly with the offer between 0 and 5 lira (e.g. for an offer of 3 Lira a manager had a chance of 60 percent of getting the information). Independent of whether a manager had obtained the information, s/he was required to divide the bonus pool over the two employees.

In total, we conducted twelve sessions with eight rounds each. Accordingly, each participant engaged in eight separate games. New groups of three participants were created at the start of every round using a stranger design matching pattern. While participants could be re-matched with another participant up to two times, they were never in a particular triad more than once, and because they interacted anonymously they did not know when, or with whom they were re-matched. After round four, all participants who played the role of manager in the first four rounds changed roles to act as employee in the second four rounds. Also, half of the participants acting as employee in the first half of the session changed roles and acted as manager in the second half.

Participants and experimental procedures

The participants are undergraduate students from a business school in The Netherlands. In total, 126 students participated in the study. The mean age of the participants was 20.4 years, with the youngest participant being 18 and the oldest 29. There were 42 (33 percent) female

participants and 84 (67 percent) males. All twelve sessions were run with either twelve or nine participants.[11] In total there were 84 different managers and 336 unique triads / games.

The students self-registered as participants in response to an invitation on the university's laboratory web site. Course credits were used as a show up fee. The actual payout in Euros for the participants was determined by randomly selecting one of the eight rounds as pay round at the end of each session and converting the Lira payoff from this round to Euros. The average amount earned was 6.58 Euro with a minimum of 0.83 Euro and a maximum of 11.93 Euro.

The laboratory consisted of a central area surrounded by 12 cubicles, each with a computer, and a control room. Upon arrival the students entered a waiting room. An instructor explained the basic procedures and told the participants that they would find a detailed set of instructions in their cubicle. Next, the participants were randomly assigned to a cubicle and started reading the instructions. The instructions explained the task and provided an example. They also explained the procedures for determination of the participants' pay-off and emphasized that they would be interacting with each other and that there was no deception of any kind. To prevent negative connotations, the instructions described the roles of the players as a division manager (supervisor) and two business unit managers instead of a manager and employees.

The computer task started automatically, about 10 minutes after the participants entered the cubicles. The participants first had to answer a set of questions about the experimental task, which served as an understanding check. Participants could not continue without having given the correct answer to all questions. After all participants had successfully completed this check, the first round began. On average, the eight rounds of the experiment took the participants 25 minutes. During the experiment, the participants had the opportunity to refer to hardcopies of Table 1,

[11] The laboratory had twelve computers available and we ran sessions with nine participants only if fewer than twelve of the registered participants showed up. In total there were six sessions with nine and six sessions with twelve participants.

which were placed on their desks. After the last round, the participants completed an exit questionnaire. This questionnaire was used to gain a better understanding of the motives behind participants' decisions. Finally, after filling in the questionnaire, the randomly selected pay round was reported, and participants collected their money and left.

IV. RESULTS

Descriptive statistics

In the exit questionnaire, we asked participants three questions regarding their involvement in the session. The items and descriptive statistics are provided in Panel C of Table 4. These indicate that most participants participated seriously in the session, made their choices after some deliberation, and cared about the outcomes of their decisions.

--- Table 4 about here ---

In analyzing our data we look at the pooled data from all 336 rounds of the experiment. Tables 4, 5 and 6 contain descriptive statistics. First, Panel A of Table 4 shows that on average employees chose effort levels equal to more than half of their endowment (5.43 Lira out of 10), signaling at least some degree of trust in managers' intentions to provide them with a fair return. The data in Panel B of Table 4 and in Tables 5 and 6 provides some more detail about the effort choices. Panel B of Table 4 shows that the mean effort level was about the same for all three types of employee players (i.e., participants who acted as employee only in the first four rounds, participants who acted as employee only in the last four rounds and participants who acted as employee in all eight rounds). An ANOVA confirms that mean effort level does not differ between types of employees *($F = 0.728$, $p = 0.485$)*. Table 5 shows the observed frequencies of individual and team effort levels. In addition, Table 6 shows how the mean effort evolved during the eight rounds of the experiment. Tables 4 and 6 also provide descriptives regarding the price offers of the

managers. The mean offer is 2.08 Lira (std. dev. = 1.67). As is clear from Table 6, the mean offer is somewhat higher in the first four rounds of the session than in the last four rounds (we control for this difference in our further analyses). Panel A of Figure 2 displays the mean offer for different values of total effort. In interpreting this figure, however, it should be noted that, as shown in Table 5, some of the averages in this figure represent relatively few observations, particularly at the more extreme levels of aggregate performance.

--- Table 5 about here ---

--- Table 6 about here ---

--- Figure 2 about here ---

Hypothesis tests

Our first hypothesis (H1) predicts that managers will generally be willing to obtain costly, non-contractible information on individual contributions to joint production. The mean price offer of 2.08 is significantly higher than zero ($p < 0$ .001), providing statistical support for H1. Of course, given that the price is truncated at zero, this statistical test is not particularly insightful into the average behavior of manager participants. A closer look at the offers reveals that a select group of eleven managers (13.1 percent) exhibit the typical behavior of the "Homo Economicus," offering 0 Lira for the information in all four rounds in which they act as managers. Three other managers also kept their offers constant despite varying levels of total effort. Two of these made the maximum offer of 5 Lira in all four rounds and the other always offered 3 Lira.

H2 predicts that managers will be more willing to obtain costly, non-contractible information as aggregate performance becomes less extreme, and H3 predicts that they will be more willing to do so as aggregate performance increases. As described earlier, together, H2 and H3 suggest a non-linear relation between the aggregate performance measure and the price offers made by managers. Specifically, the price offers will increase as aggregate performance

increases, up to the point where aggregate performance (total effort) is 10. After this point the price offers will either increase at a lower rate or will decrease as aggregate performance increases from 10 to 20. Thus, we are able to test both H2 and H3 by examining the pattern of data. We do so using the following regression:

$$PRICE = \alpha + \beta_1 \cdot EXTREME + \beta_2 \cdot HIGHPERF + \beta_3 \cdot HIGHPERF \cdot EXTREME + \beta_4 \cdot EMPEXP \quad (1)$$

where *PRICE* = the price offer made by the manager,

EXTREME = the absolute value of (aggregate performance minus 10),

HIGHPERF = a dummy variable which has the value 1 if aggregate performance ≥ 10 and 0 otherwise,

HIGHPERF·EXTREME = the interaction between *HIGHPERF* and *EXTREME*, and

EMPEXP = a dummy variable which has the value of 1 if the manager has previous experience as an employee, and 0 otherwise.

We take several important steps to ensure that the statistical tests provided by this regression are appropriate for the experimental design used. First, we exclude observations for which aggregate performance = zero or 20 (*EXTREME* = 0), allowing for a discontinuity at these extreme observations.[12] Next, because we collect four different observations of the price offer from each participant acting as a manager, our data violate the assumption of independence. To correct for this violation, we calculate robust estimators (also known as Huber-White or sandwich estimators), using the Generalized Estimating Equations (GEE) module of SPSS. This method provides estimates that are corrected for cluster-correlated data such as ours (Wooldridge 2003,

[12] Note that if total effort is either zero or 20, there is no uncertainty about the individual effort levels and consequently we would not expect the manager to make a positive price offer for the information (i.e. our theory only applies for values of *EXTREME* > 0). Thus, including observations with *EXTREME* = 0 in our analysis might obscure an existing linear relationship between *EXTREME* and *PRICE*. In fact, our results are inferentially identical if we do include these observations, albeit the coefficient on the interaction term drops to marginal significance (p = 0.054). If we use an independent correlation structure instead of an auto-regressive correlation structure, the interaction term is significant at the 0.05 level even if we include cases in which *EXTREME* = 0.

2006).[13] Finally, because the participants changed roles after four rounds, we include a dummy variable (*EMPEXP*) to control for any effects of prior experience as an employee on decision making as a manager.

Table 7 and Figure 3 present the regression results. Importantly, the x-axis on the graph in Figure 3 is aggregate performance, whereas the regression itself does not include a term for aggregate performance. Therefore, it is useful to translate the regression output to its graphical representation.[14] When aggregate performance equals 10, the most moderate level of aggregate performance, the variable *EXTREME* equals 0. Therefore, the intercept α equals the price expected as aggregate performance approaches 10 from the left. This intercept plus β_2, the coefficient on the dummy variable *HIGHPERF*, equals the price expected as aggregate performance approaches 10 from the right. While these two points are separated in the figure, it should be noted that the difference (i.e., β_2) is not significant at conventional levels *($p = 0.118$).* The negative of β_1, the coefficient on *EXTREME*, equals the slope of the regression line to the left of 10, whereas the sum of β_1 and β_3, the coefficient on the interaction term, equals the slope of the regression line to the right of 10.

--- Table 7 about here ---

--- Figure 3 about here ---

The pattern of data presented in Figure 3 is quite similar to that predicted in Figure 1, Panel A. Specifically, the price offers increase as aggregate performance increases, up to the point

[13] We run GEE with a fixed scale value of 1 and with an auto-regressive correlation structure. The auto-regressive correlation structure allows us to take into account that observations which are closer together in time (e.g. a manager's first and second price offer) are likely to have a stronger correlation than observations that are further away from each other (e.g. a manager's first and fourth price offer) (e.g. Zorn, 2001). Our findings are qualitatively similar if we use an independent correlation structure (assume no correlation between the different observations for a specific participant). In that case the interaction coefficient also remains significant if we include the observations for which *EXTREME* = 0.

[14] Note that for expositional clarity, in Figure 1, Panel B, we also recast the predicted pattern, using *EXTREME* rather than aggregate performance along the x-axis.

where aggregate performance (total effort) is 10. At this point, there is a downward shift in the line, but this shift creates an asymmetry, such that the slope of the line to the left of 10 is much steeper than that to the right of 10.

We verify that the pattern is statistically significant. First, the significantly negative β_1 *(p = 0.001)* suggests that as aggregate performance approaches 10 from the left, price offers increase. We next provide statistical evidence of the fairness effect (H2), by demonstrating that there is a kink in the line at the point where aggregate performance = 10. Recall that the slope to the left of 10 equals $-\beta_1$ and that the slope to the right of 10 equals $\beta_1 + \beta_3$. Therefore, to test for a slope change at 10, we test a planned comparison that $\beta_1 + \beta_3 < -\beta_1$. This test provides evidence of a statistically significant *(p = 0.014)* reduction in the slope at the point where aggregate performance equals 10. Finally, the significantly positive β_3 *(p = 0.018)* indicates an asymmetry in the relation between the aggregate performance measure and the price offer, which, as described earlier, provides evidence of the trust reciprocity effect predicted in H3.

The H3 result is notable, because it works in the opposite direction of the effect described in the recent analytic literature. Rajan and Reichelstein (2009) demonstrate that under certain assumptions, non-contractible performance information will be less valuable for the allocation of a discretionary bonus pool when contractible measure outcomes are higher. We do not interpret our results as conflicting with those of Rajan and Reichelstein's (2009), but rather as complementing them. Rajan and Reichelstein's (2009) results are based on theories of optimal contracting under traditionally modeled utility functions. Further, they do not model the joint production setting, which is central to our investigation. Our results would seem to create a boundary condition for their results, suggesting that in the team setting, managers' social preferences could mitigate Rajan and Reichelstein's (2009) effect, such that managers would place greater value on non-contractible information when the outcomes of objective performance metrics are higher.

In summary, we find that managers' willingness to obtain costly, non-contractible information that allows them to *ex post* distinguish the individual efforts of their employees follows the predicted patterns. Over the whole range of observed total effort levels, managers are prepared to obtain costly, non-contractible information. However, their willingness to do so depends on the outcome of the contractible performance measure, aggregate performance. Specifically, the willingness to obtain costly, non-contractible information increases as aggregate performance becomes more extreme, which is consistent with the predicted fairness effect. Further, the willingness to obtain this information increases with aggregate performance, providing support for the predicted trust reciprocity effect. Because we find that for relatively high levels of total effort, the price offer is a decreasing function of aggregate performance, our results suggest that in our sample the fairness effect was stronger than the reciprocity effect.

Supplemental analysis

This section contains some additional analysis of our results. First, we use data from the exit questionnaire to gain a better understanding of the decisions made by the manager participants. Next, we examine and discuss the allocation decisions for managers who did and did not obtain information on individual contributions to joint production.

To provide additional evidence on the processes responsible for the effects we observe, we extract a subset of manager participants who report having placed more weight on fairness than on trust reciprocity when making their decisions, and then compare the pattern of their responses to what our theory would predict for managers driven solely by fairness concerns, with trust reciprocity playing no role. Similarly, we extract a subset of manager participants who report having placed more weight on trust reciprocity than on fairness concerns when making their

decisions, and then compare the pattern of their responses to what our theory would predict for managers driven solely by trust reciprocity, with fairness concerns playing no role.

We extract our subsets of managers using the results of the exit questionnaire. This questionnaire contained a set of items that allow us to evaluate the motivations behind managers' price offers and allocation decisions. Table 8 provides descriptive statistics on these questionnaire items. Note that items 3 through 7 relate to the manager's motivation to establish a social norm of fair cooperation (ensure that cooperators were rewarded and slackers were punished). Similarly, items 8 through 11 relate to manager's desire to reciprocate the trust that investing employees placed in them (items 8 to 11). We subtract the sum of items 3 to 7 from the sum of items 8 to 11, resulting in a difference score, representing the relative importance of trust reciprocity vs. fairness to the manager. Those managers who score in the bottom 25% are labeled HIGHFAIR managers, and those who score in the top 25% are labeled HIGHTRUST managers.

--- Table 8 about here ---

Recall from footnote 9 that if managers' decisions were driven solely by fairness concerns and that trust reciprocity played no role, we would expect to see a perfectly symmetric relation between the aggregate performance measure and the managers' price offers, with the offers increasing up to an aggregate performance measure of 10 and then decreasing at the same rate beyond that point. We re-run our primary regression analysis (Equation 1), using only the data for HIGHFAIR managers, and graph the results in Figure 4, similar to the presentation of the main results in Figure 3. The results of this analysis must be interpreted with some caution, because they are based on only 25% of our manager participants and thus represent a small sample size. Further, we are comparing them to a benchmark based on managers who place zero weight on trust reciprocity, whereas these managers simply place less weight on trust reciprocity than on

fairness concerns. Nonetheless, the visual pattern of results is quite similar to the pattern we would expect for managers for whom fairness, but not trust reciprocity, is a primary concern.

--- Figure 4 about here ---

Recall further that if managers' decisions were driven solely by trust reciprocity and that fairness played no role, we would expect no change in the slope at 10. Rather, we would expect a consistent and monotonic increase in price offers as the aggregate performance measure increased from 0 to 20. We re-run our primary regression analysis (Equation 1), using only the data for HIGHTRUST managers, and graph the results in Figure 5, similar to the presentation of the main results in Figure 3. Using similar reasoning to that above, the results of this analysis must be interpreted with some caution, because they represent a small sample size and because, we are comparing them to a benchmark based on managers who place zero weight on fairness, whereas these managers simply place less weight on fairness concerns than on trust reciprocity. Nonetheless, the visual pattern of results is quite similar to the pattern we would expect for managers for whom trust reciprocity, but not fairness, is a major concern.

--- Figure 5 about here ---

We compare these results statistically by running a regression analysis using both the HIGHFAIR and the HIGHTRUST managers for the cases in which the aggregate measure was moderate to high (i.e., *HIGHPERF* = 1), since only in this region the fairness and trust argument lead to different predictions. We test the following regression model:

$$PRICE = \alpha + \beta_1 \cdot EXTREME + \beta_2 \cdot TRUSTvsFAIR + \beta_3 \cdot EXTREME \cdot TRUSTvsFAIR + \beta_4 \cdot EMPEXP, \quad (2)$$

where *PRICE* = the price offer made by the manager,

EXTREME = the absolute value of (aggregate performance minus 10),

TRUSTvsFAIR = a dummy variable which has the value 1 if the participant is a HIGHTRUST manager and the value 0 if the participant is a HIGHFAIR manager

EXTREME·TRUSTvsFAIR = the interaction between *EXTREME* and *TRUSTvsFAIR,*and

EMPEXP = a dummy variable which has the value of 1 if the manager has previous experience as an employee, and 0 otherwise.

A significantly positive coefficient on the interaction would suggest that the slope to the right of 10 is more positive for the HIGHTRUST than HIGHFAIR managers. As presented in Table 9, the interaction is significant *(p = 0.014)*. Thus, we conclude that the patterns represented in Figures 4 and 5 are significantly different. More importantly, they are different in precisely the way that our theory would predict. That is, managers who report to care more about trust reciprocity than fairness increase their price offers as the aggregate performance measure increases from 10 to 20, whereas those who report to care more about fairness than trust reciprocity decrease their price offers over this same range. Thus, this analysis provides corroborating evidence that our main results are driven by the social preferences addressed by our theory.

--- Table 9 about here ---

Regarding the managers' allocation decisions, Table 10 compares the allocation decisions made by managers who obtained the costly, non-contractible information on individual contributions to joint production to the decisions of those who did not. The price offers of the managers were high enough to obtain the information in 131 (39 percent) of the triads and not high enough in 205 (61 percent) of the triads. Managers that obtained the information used it to allocate the bonus pool in proportion to individual effort in the vast majority of cases (111 cases, 85 percent). The data in Table 8 also show that in most cases (130, 63 percent) managers who allocated the bonus pool without the information chose a 50-50 split. In the remaining 75 cases the allocations range from 0-100 splits to 49-51 splits. Closer analysis of the data shows an interesting

pattern. In the 89 cases where managers offered zero Lira for the information, and accordingly did not obtain it, the employees were both allocated exactly half of the bonus pool in 82 percent of the times (73 cases). On the other hand, in the 116 cases where the offer was positive but not high enough to obtain the information, the bonus is split equally in only 49 percent of the situations (57 cases). Our theory does not enable us do draw conclusions about the causal mechanisms responsible for this pattern. However, it is tempting to speculate that managers who are unwilling to incur a cost to obtain individual effort information compensate this selfish choice by at least acting as a "predictable" manager. Managers whose price offers are positive but too low to receive the information, on the other hand, may feel frustrated that their unselfish behavior does not pay off, and dividing the bonus pool in unequal random portions may be a way to relieve this frustration.

--- Table 10 about here ---

V. DISCUSSION AND CONCLUSIONS

This study uses an experiment to examine the effect of social preferences on managers' allocations of discretionary bonus pools. We find that very few managers behave in accordance with the traditional model of "economic man." Most managers are willing to obtain costly, non-contractible information that allows them to unravel the contractible metric of aggregate performance, and thus to reward their employees' relative efforts. However, this willingness is influenced by the outcome of aggregate performance. Specifically, the willingness increases as aggregate performance becomes less extreme, and (holding extremeness constant) as aggregate performance increases. These results are consistent with models that incorporate social preferences and, in particular, preferences for fairness and trust reciprocity. Our findings have implications for both the accounting literature and research on human altruism and third-party intervention.

The study contributes to accounting research, as it is one of the first to systematically investigate how non-selfish motivations influence the use of non-contractible information in discretionary bonus allocations. While the existing literature has shown that relational contracting and discretionary bonus pools may be efficient mechanisms to solve managerial incentive problems, this literature has not taken into account that information about the relative effort of agents is not freely available, but can be obtained at some cost to the manager. This may be problematic because managers do not always have monetary incentives to use their discretion in a value-enhancing way. Our study shows that even in the absence of such monetary incentives, social preferences can motivate managers to obtain costly, non-contractible information in order to evaluate and reward their employees according to their efforts. This may help explain why discretionary performance evaluation is so common in practice, even in situations where managers do not have explicit incentives to use their discretion to protect their employees from the inherently incomplete nature of objective performance measures.

Future research is needed to shed more light on how self-interest, cognitive biases and social preferences interact to influence the use of discretion. For example, researchers could examine how incentive structures of both employees and managers influence evaluation outcomes. They could also investigate how evaluations are affected by individual differences and by manager-employee relationship characteristics such as a history of working together. Another potentially fruitful direction for future research is examination of how social preferences influence incentive contracting and other accounting-related issues such as budgeting processes and transfer pricing mechanisms. Our research shows how theories and research methods originating in behavioral and experimental economics can inform accounting research in this respect.

Our results are also of interest to researchers in economics and related fields, as this study is one of the earliest to focus on a third party's decisions regarding the allocation of resources

between two other parties, and the first to deal with third parties' intervention in situations of *potential* unfairness. Our research suggests third parties do have preferences regarding such allocations and are willing to incur a cost to ensure that resources are allocated according to these preferences. The findings furthermore show that paying for fairness is not a dichotomous choice, as willingness to pay increases with the potential unfairness of the situation. Thus, individuals seem to make a deliberate trade-off between the costs and benefits of self-sacrificing acts of fairness. Our study also indicates that managers' willingness to pay is not solely driven by inequity aversion (c.f. Bolton and Ockenfels 1993; Fehr and Schmidt 1999)[15]. Instead, it seems to originate in an experienced need to enforce a social norm of fair cooperation and to repay trust (Berg *et al*. 1995; Fehr and Fischbacher 2004b). Consequently, the results presented in this paper are supportive of the notion of strong reciprocity (Fehr and Fischbacher 2004b; Fehr *et al*. 2002).

In conclusion, the findings in this paper highlight the need for researchers in economics and related fields to continue to develop, refine and test models of apparently unselfish human behavior in economic situations. In particular, research on third party observation of – and intervention in – economic transactions between other individuals is likely to be fruitful, as this important issue is currently understudied. For example, future studies should help us understand why in our study managers who are prepared to pay for fairness are more likely to provide completely random, and therefore potentially unfair, rewards and punishments if their price offer falls short and their initial investment in fairness does not pay off. Such insights in the fundamental drivers of people's economic decision-making behavior should also lead to a better understanding of accounting phenomena.

[15] To illustrate, in 26.0 percent of the cases in which the manager obtains the information about the individual effort levels, the bonus pool is allocated in such a way that the difference between the final payoffs of the employees is larger than the difference that would have resulted had the pool been split equally. Also, in 37.4 percent of the cases, managers use the information to provide at least one employee with a payoff that is higher than the manager's own payoff.

REFERENCES

Bailey, W. J., G. Hecht, and K. L. Towry. 2009. Dividing the pie: The influence of managerial discretion extent on bonus pool allocation. Working Paper.

Baiman, S. 1982. Agency research in managerial accounting: a survey. *Journal of Accounting Literature* 1: 154-213.

———, and M. V. Rajan. 1995. The informational advantages of discretionary bonus schemes. *The Accounting Review* 70 (4): 557-579.

Baker, G., R. Gibbons, and K. J. Murphy. 1994. Subjective performance measures in optimal incentive contracts. *Quarterly Journal of Economics* 109: 1125-1156.

———, ———, and ———. 2002. Relational contracts and the theory of the firm. *Quarterly Journal of Economics* 117 (1): 39-84.

Becker, G. M., M. H. DeGroot, and J. Marschak, 1964. Measuring utility by a single-response sequential method. *Behavioral Science* 9(2): 226-232.

Bendor, J., and P. Swistak. 2001. The evolution of norms. *American Journal of Sociology* 106 (6): 1493-1545.

Berg, J. E., J. W. Dickhaut, and K. McCabe. 1995. Trust, reciprocity, and social history. *Games and Economic Behavior* 10: 122-142.

Bohm, P., J. Lindén, and J. Sonnegård. 1997. Eliciting Reservation Prices: Becker-DeGroot-Marschak Mechanisms vs. Markets. *The Economic Journal* 107 (443): 1079-1089.

Bol., J. 2009a. Subjectivity in compensation contracting. *Journal of Accounting Literature*. Forthcoming.

———. 2009b. The determinants and performance effects of supervisor bias. Working Paper.

——— and S. D. Smith. 2009. Spillover effects in subjective performance evaluation: bias, fairness, and controllability. Working Paper.

Bolton, G. E., and A. Ockenfels. 1993. ERC: A theory of equity, reciprocity, and competition. *American Economic Review* 90 (1): 166-193.

Brüggen, A., and M. Strobel. 2007. Real effort versus chosen effort in experiments. *Economics Letters* 96 (2): 232-236.

Callahan, C. M., E. A. Gabriel and B. J. Sainty. 2006. A review and classification of experimental economics research in accounting. *Journal of Accounting Literature* 25: 59-126.

Camerer, C. 2003. *Behavioral Game Theory: experiments in strategic interaction*. Princeton NJ: Princeton University Press.

Carpenter, J. P., P. H. Matthews, and O. Ong'ong'a. 2004. Why punish? Social reciprocity and the enforcement of prosocial norms. *Journal of Evolutionary Economics* 14 (4): 407-429.

Charness, G., and M. Rabin. 2002. Understanding social preferences with simple tests *Quarterly Journal of Economics* 117 (3): 817-869.

Cropanzano, R., Z. S. Byrne, D. R. Bobocel, and D. E. Rupp. 2001. Moral virtues, fairness heuristics, social entities, and other denizens of organizational justice. *Journal of Vocational Behavior* 58: 164-209.

De Quervain, J., U. Fischbacher, V. Treyer, M. Schellhammer, U. Schnyder, A. Buck, and E. Fehr. 2004. The neural basis of altruistic punishment. *Science* 425: 785-791.

Drake, A. R., and S. F. Haka. 2008. Does ABC information exacerbate hold-up problems in buyer-supplier negotiations? *The Accounting Review* 83 (1): 29-60.

Dufwenberg, M., and U. Gneezy. 2000. Measuring beliefs in an experimental lost wallet game. *Games and Economic Behavior* 30: 163-182.

Ederhof, M. 2008. Discretion in bonus plans. Working Paper.

Evans, J. H., R. L. Hannan, R. Krishnan, and D. V. Moser. 2001. Honesty in managerial reporting. *The Accounting Review* 76 (4): 537-559.

Fehr, E., and K. M. Schmidt. 1999. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114: 817-851.

———, and ———. 2003. Theories of fairness and reciprocity - evidence and economic applications. In *Advances in Economic Theory: Eighth World Congress of the Economic Society*, edited by M. Dewatripont, L. P. Hansen and S. J. Turnovsky. Cambridge, UK: Cambridge University Press.

———, and U. Fischbacher. 2002. Why social preferences matter - The impact of non-selfish motives on competition, cooperation and incentives. *The Economic Journal* 112 (478): C1 – C33.

———, and ———. 2004a. Social norms and human cooperation. *Trends In Cognitive Sciences* 8 (4): 185-190.

———, and ———. 2004b. Third-party punishment and social norms. *Evolution And Human Behavior* 25 (2): 63-87.

———, ———, and S. Gaechter. 2002. Strong reciprocity, human cooperation and the enforcement of social norms. *Human Nature* 13: 1-25.

Fischbacher, U. 2007. Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10 (2): 171-178.

Fisher, J. G., L. A. Maines, S. A. Peffer, and G. B. Sprinkle. 2005. An experimental investigation of employer discretion in employee performance evaluation and compensation. *The Accounting Review* 80 (2): 563-583.

Fiske, A. P. 1991. The cultural relativity of selfish individualism: anthropological evidence that humans are inherently social. In *Prosocial Behavior*, edited by M. Clark. Newbury Park, CA: Sage.

Gibbs, M., K. A. Merchant, W. A. van der Stede, and M. E. Vargus. 2004. Determinants and effects of subjectivity in incentives. *The Accounting Review* 79 (2): 409-436.

Gintis, H. 2000. Strong reciprocity and human sociality. *Journal of Theoretical Biology* 206: 169-179.

Hannan, R. L. 2005. The combined effect of wages and firm profit on employee effort. *The Accounting Review* 80 (1): 167-188.

_____, F. W. Rankin, and K. L. Towry. 2009. Flattening the organization: the effect of organizational reporting structure on honesty in managerial reporting. *Working Paper.*

Holmström, B. 1979. Moral hazard and observability. *Bell Journal of Economics* 10: 74-91.

Ittner, C. D., D. F. Larcker, and M. W. Meyer. 2003. Subjectivity and the weighting of performance measures: evidence from a balanced scorecard. *The Accounting Review* 78 (3): 725-758.

Kachelmeier, S. J., and K. L. Towry. 2002. Negotiated transfer pricing: is fairness easier said than done? *The Accounting Review* 77 (3): 571-593.

Kahneman, D., J. L. Knetsch, and R. Thaler. 1986. Fairness as a constraint on profit seeking: entitlements in the market. *American Economic Review* 76 (4): 728-741.

Lambert, R. 1985. Variance investigation in agency settings. *Journal of Accounting Research* 23: 633-647.

Lipe, M. G., and S. E. Salterio. 2000. The balanced scorecard: judgmental effects of common and unique performance measures. *The Accounting Review* 75 (3): 283-298.

MacLeod, W. B. 2003. Optimal contracting with subjective evaluation. *American Economic Review* 93 (1): 216-240.

Matsumura, E. M., and J. Y. Shin. 2006. An empirical analysis of an incentive plan with relative performance measures: Evidence from a postal service. *The Accounting Review* 81 (3): 533-566.

Moers, F. 2005. Discretion and bias in performance evaluation: the impact of diversity and subjectivity. *Accounting Organizations and Society* 30 (1): 67-80.

Mohtashemi, M., and L. Mui. 2003. Evolution of indirect reciprocity by social information: the role of trust and reputation in evolution of altruism. *Journal of Theoretical Biology* 223: 523-531.

Prendergast, C., and R. Topel. 1993. Discretion and bias in performance evaluation. *European Economic Review* 37: 355-365.

Rajan, M. V. and S. Reichelstein. 2006. Subjective performance indicators and discretionary bonus pools. *Journal of Accounting Research* 44 (3): 585-618.

Rajan, M. V. and S. Reichelstein. 2009. Objective versus subjective indicators of managerial performance. *The Accounting Review* 84(1): 209-237.

Reuben, E. and F. van Winden (2008). Social ties and coordination on negative reciprocity: The role of affect. *Journal of Public Economics* 92: 34-53.

Shadish, W. R., T. D. Cook, and D. T. Campbell. 2002. *Experimental and quasi-experimental designs for generalized causal inference*. Boston & New York: Houghton Mifflin Company.

Turillo, C. J., R. Folger, J. J. Lavelle, E. E. Umphress, and J. O. Gee. 2002. Is virtue its own reward? Self-sacrificial decisions for the sake of fairness. *Organizational Behavior and Human Decision Processes* 89 (1): 839-865.

Van Winden, F. (2007). Affect and fairness in economics. *Social Justice Research* 20 (1): 35-52.

Wooldridge, J. M. 2003. Cluster-sample methods in applied econometrics. *American Economic Review* 93 (2): 133-138.

———. 2006. Cluster-sample methods in applied econometrics: An extended analysis. Working Paper, Michigan State University.

Young, R. 1986. A note on: Economically optimal performance evaluation and control systems: The optimality of two-tailed investigations. *Journal of Accounting Research* 24: 231-240.

Zorn, C. J. W. 2001. Generalized Estimation Equations for correlated data: A review with applications. *American Journal of Political Science*, 45 (2): 470-490.

TABLE 1

Bonus Pool Size as a Function of Employee Effort [a]

Employee B Effort	Employee A Effort: 0	1	2	3	4	5	6	7	8	9	10
0	0.00	1.50	3.00	4.50	6.00	7.50	9.00	10.50	12.00	13.50	15.00
1	1.50	3.00	4.50	6.00	7.50	9.00	10.50	12.00	13.50	15.00	16.50
2	3.00	4.50	6.00	7.50	9.00	10.50	12.00	13.50	15.00	16.50	18.00
3	4.50	6.00	7.50	9.00	10.50	12.00	13.50	15.00	16.50	18.00	19.50
4	6.00	7.50	9.00	10.50	12.00	13.50	15.00	16.50	18.00	19.50	21.00
5	7.50	9.00	10.50	12.00	13.50	15.00	16.50	18.00	19.50	21.00	22.50
6	9.00	10.50	12.00	13.50	15.00	16.50	18.00	19.50	21.00	22.50	24.00
7	10.50	12.00	13.50	15.00	16.50	18.00	19.50	21.00	22.50	24.00	25.50
8	12.00	13.50	15.00	16.50	18.00	19.50	21.00	22.50	24.00	25.50	27.00
9	13.50	15.00	16.50	18.00	19.50	21.00	22.50	24.00	25.50	27.00	28.50
10	15.00	16.50	18.00	19.50	21.00	22.50	24.00	25.50	27.00	28.50	30.00

[a] We use an abstract notion of effort as an incurred cost. For example, an employee with an effort level of 5 incurs a cost of 5 Lira (the experimental currency).

TABLE 2
Experimental Variables

Variable name	Code	Value
Employee A endowment	N_a	10 Lira
Employee B endowment	N_b	10 Lira
Manager endowment	N_s	15 Lira
Employee A effort	E_a	$0 \leq E_a \leq N_a$
Employee B effort	E_b	$0 \leq E_b \leq N_b$
Aggregate performance measure	T	$E_a + E_b$
Bonus pool	B	1.5 * T
Price offer of manager	P_{offer}	0 Lira $\leq P_{offer} \leq$ 5 Lira
Actual price of information	P_{act}	0 Lira $\leq P_{act} \leq$ 5 Lira (determined at random)
Bonus allocated to employee A	B_a	$B - B_b$
Bonus allocated to employee B	B_b	$B - B_a$
Employee A final pay-off	F_a	$N_a - E_a + B_a$
Employee B final pay-off	F_b	$N_b - E_b + B_b$
Manager final pay-off	F_s	$N_s - P_{act}$ for $P_{offer} \geq P_{act}$ and N_s for $P_{offer} < P_{act}$

TABLE 3:
Session Time Line per Round

Stage	
1	New groups are formed (one manager and two employees)
2	The employees and the manager receive their endowments N_a, N_b and N_s.
3	The employees decide on their effort levels (i.e., monetary costs, E_a and E_b, up to the amount of the initial endowments, N_a and N_b).
4	All participants observe the aggregate performance measure $T = (E_a + E_b)$ and the available bonus pool B.
5	The manager indicates the maximum price P_{offer} s/he is prepared to pay to learn E_a and E_b.
6	The computer randomly determines the actual price P_{act}.
7	If $P_{offer} \geq P_{act}$, the manager learns about E_a and E_b and his/her pay-off is reduced with P_{act}.
8	The manager decides how to allocate the bonus pool B over the two employees (i.e. decides about B_a and B_b).
9	The employees receive their allocations of the bonus pool.
10	All participants learn about the round's pay-offs, F_a, F_b and F_s, as well as the offer of the manager P_{offer} and the actual price P_{act}.

TABLE 4:
Descriptive Statistics

Panel A: Primary Variables

	N	Min	Max	Median	Mean	SD
Employee effort	672	0	10	6	5.43	3.26
Manager Price Offer	336	0	5	2	2.08	1.67
Employee payoff	672	0	28	12.99	12.71	3.14
Manager payoff	336	10.84	15	15	14.31	1.10

Panel B: Employees' Mean Effort Over Four or Eight Rounds as Employee

Player type	N	Mean	SD
Employee in first half	42	5.11	2.79
Employee in second half	42	5.17	2.57
Employee in both halves	42	5.72	2.22
Total	126	5.33	2.53

Panel C: Post-Experimental Items Measuring Participants' Involvement

Item	Mean	SD	Theoretical range	Actual range
I participated seriously in this study	4.68	0.57	1 – 5	2 – 5
I thought about my choices before making any decisions	4.43	0.71	1 – 5	2 – 5
I really cared about the outcomes of my decisions	4.52	0.76	1 – 5	1 – 5

TABLE 5
Observed Frequencies Individual and Team Effort

Individual effort

Effort	Frequency	Percentage	Cumulative Percentage
0	87	12.9	12.9
1	24	3.6	16.5
2	36	5.4	21.9
3	39	5.8	27.7
4	65	9.7	37.4
5	78	11.6	49.0
6	74	11.0	60.0
7	57	8.5	68.5
8	75	11.2	79.6
9	28	4.2	83.8
10	109	16.2	100.0
Total	672	100.0	

Aggregate Performance Measure (Team Effort)

Team Effort	Frequency	Percentage	Cumulative Percentage
0	3	0.9	0.9
1	5	1.5	2.4
2	9	2.7	5.1
3	3	0.9	6.0
4	8	2.4	8.3
5	11	3.3	11.6
6	11	3.3	14.9
7	18	5.4	20.2
8	34	10.1	30.4
9	25	7.4	37.8
10	32	9.5	47.3
11	35	10.4	57.7
12	19	5.7	63.4
13	22	6.5	69.9
14	21	6.3	76.2
15	24	7.1	83.3
16	24	7.1	90.5
17	12	3.6	94.0
18	9	2.7	96.7
19	3	0.9	97.6
20	8	2.4	100.0
Total	336	100.0	

TABLE 6
Mean Employee Effort Levels and Manager Price Offers Across Rounds

Employee Effort

	Round								
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**Overall**
Mean	5.52	5.52	5.68	5.35	5.14	5.62	5.57	5.02	5.43
SD	2.75	3.17	3.28	3.20	3.16	3.28	3.49	3.73	3.26
N	84	84	84	84	84	84	84	84	672

Manager Price Offers

	Round								
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**Overall**
Mean	2.35	2.46	2.27	2.39	1.66	1.81	1.95	1.74	2.08
SD	1.43	1.57	1.75	1.62	1.54	1.64	1.89	1.79	1.67
N	42	42	42	42	42	42	42	42	336

TABLE 7
Results of GEE Analysis[a]

$PRICE = \alpha + \beta_1 \cdot EXTREME + \beta_2 \cdot HIGHPERF + \beta_3 \cdot HIGHPERF \cdot EXTREME + \beta_4 \cdot EMPEXP$

	B	Std. Error	Wald Chi Square	df	p-value (two-tailed)
Intercept	1.897	0.252	132.556	1	<0.001
EXTREME	-0.177	0.053	10.954	1	0.001
HIGHPERF	-0.386	0.247	2.444	1	0.118
*HIGHPERF*EXTREME*	0.151	0.064	5.622	1	0.018
EMPEXP	-0.526	0.289	3.317	1	0.069

PRICE = the price offer made by the manager
EXTREME = the absolute value of (aggregate performance minus 10).
HIGHPERF = a dummy variable which has the value 1 if aggregate performance ≥ 10 and 0 otherwise.
EMPEXP = a dummy variable which has the value of 1 if the manager has previous experience as an employee, and 0 otherwise.

Planned Contrast: $\beta_1 + \beta_3 < -\beta 1$ Wald Chi Square = 6.10, p = 0.014

[a] To correct for cluster-correlated data (due to manager participants each providing four observations), we calculate robust estimators (also known as Huber-White or sandwich estimators), using the Generalized Estimating Equations (GEE) module of SPSS. We run GEE with a fixed scale value of 1 with an auto-regressive correlation structure. The auto-regressive correlation structure allows us to take into account that observations which are closer together in time (e.g., a manager's first and fourth price offer) are likely to have a stronger correlation than observations that are further away from each other.

TABLE 8
Descriptive Statistics - Exit Questionnaire Items[a]

Item		Mean	SD	t	p
(1)	In general I was curious about the separate investments of the two BU managers	3.86	1.28	6.13	0.000
(2)	In general I wanted to know whether one BU manager had invested more than the other	3.89	1.28	6.39	0.000
(3)	I wanted to reward BU managers who acted cooperatively	3.87	1.30	6.14	0.000
(4)	I wanted to punish BU managers who did not act cooperatively	3.61	1.38	4.03	0.000
(5)	I thought it was important that BU managers who tried to get more than their fair share got punished	3.23	1.27	1.63	0.107
(6)	I thought it was important that BU managers who acted in the common interest got rewarded	3.70	1.08	5.94	0.000
(7)	I thought it was important that BU managers who acted in the common interest got at least a fair return	3.99	1.05	8.65	0.000
(8)	I wanted to repay the trust that BU managers placed in me by investing part of their base amount	3.58	1.12	4.77	0.000
(9)	I wanted to reward those BU managers who expected me to be a fair manager	3.71	1.14	5.76	0.000
(10)	I wanted to punish those BU managers who did not expect me to be a fair manager	2.92	1.08	-0.71	0.481
(11)	I did not want to disappoint BU managers who trusted me to reward high investments	3.43	1.26	3.11	0.003

[a] All items scored on a five point Likert scale (fully disagree – fully agree). Answers on all items cover the whole theoretical range (1 to 5). The last two columns give the t-statistic and associated two-sided p value for a t-test if the mean score differs from the theoretical mean of 3.

TABLE 9
Results of GEE Analysis for HIGHPERF = 1 subsample, including only HIGHTRUST and HIGHFAIR Managers[a]

$PRICE = \alpha + \beta_1 \cdot EXTREME + \beta_2 \cdot TRUSTvsFAIR + \beta_3 \cdot EXTREME \cdot TRUSTvsFAIR + \beta_4 \cdot EMPEXP$

	B	Std. Error	Wald Chi Square	df	p-value (two-tailed)
Intercept	2.403	0.459	27.364	1	<0.001
EXTREME	-0.126	0.090	1.977	1	0.160
TRUSTvsFAIR	-0.757	0.554	1.868	1	0.172
*EXTREME*TRUSTvsFAIR*	0.299	0.121	6.080	1	0.014
EMPEXP	-0.229	0.427	0.289	1	0.591

PRICE = the price offer made by the manager
EXTREME = the absolute value of (aggregate performance minus 10).
TRUSTvsFAIR = a dummy variable which has the value of 1 if the manager is a HIGHTRUST manager and 0 if the manager is a HIGHFAIR manager.
EMPEXP = a dummy variable which has the value of 1 if the manager has previous experience as an employee, and 0 otherwise.

[a] To correct for cluster-correlated data (due to manager participants each providing four observations), we calculate robust estimators (also known as Huber-White or sandwich estimators), using the Generalized Estimating Equations (GEE) module of SPSS. We run GEE with a fixed scale value of 1 with an auto-regressive correlation structure. The auto-regressive correlation structure allows us to take into account that observations which are closer together in time (e.g., a manager's first and fourth price offer) are likely to have a stronger correlation than observations that are further away from each other.

TABLE 10
Allocation Decisions of Managers

Offer		**Obtained Information?**		**50/50 allocation when information was not obtained?**		**Allocation in proportion to relative effort?**[a]	
Zero	89	No	89	No	16		
				Yes	73		
				Total	89		
Positive	247	No	116	No	59		
				Yes	57		
				Total	116		
		Total No	205	Total No	75		
				Total Yes	130		
		Yes	131			No	20
						Yes	111
Total	336		336		205		131

[a] Considered in proportion to relative effort if proportions of bonus pool allocations are within 10 percentage points of relative efforts. We chose this measure because some managers round their allocation to the nearest 10 percent (e.g. if the information shows that A has provided 33 percent of total effort and B has provided 67 percent the manager chooses a 30-70 split).

FIGURE 1
Expected Results: Willingness to Incur Cost as a Function of Aggregate Performance Measure

Panel A: Aggregate Performance Measure (Total Effort) on x-axis[a]

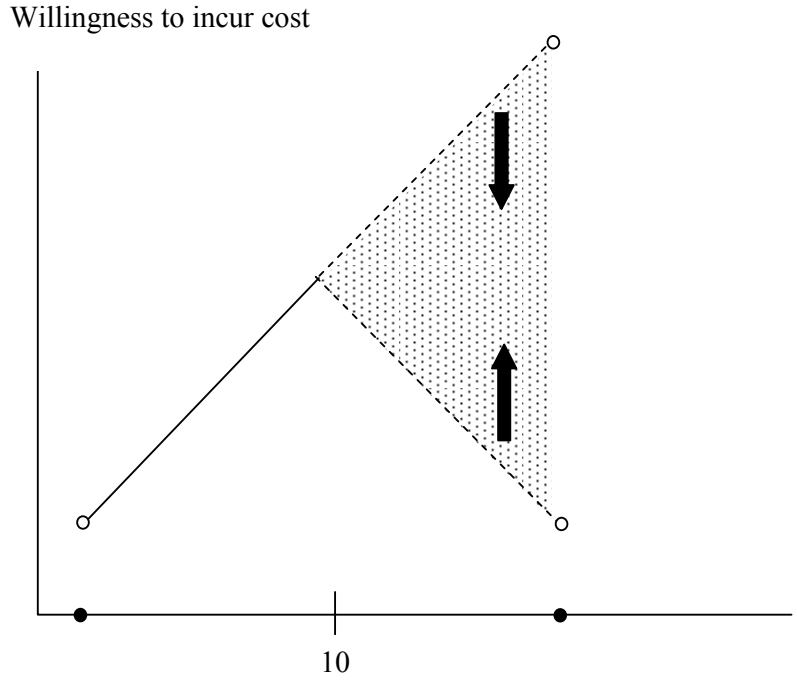

[a] Note that as the aggregate measure increases from low levels of effort to moderate levels (i.e., 10), the slope of the line representing the manager's willingness to pay for information is positive, reflecting both the fairness (H2) and the trust reciprocity (H3) effects. As the aggregate measure then increases from moderate to high levels, the slope of this line becomes less steep, reflecting the positive effect of trust reciprocity (H3) and the offsetting negative effect of fairness (H2). Depending on the relative sizes of these effects (which we do not predict), the slope may remain positive or become negative. Importantly, if the slope becomes negative (indicating a stronger effect for fairness than for trust reciprocity), we would expect this negative slope to be less steep than the positive slope to the left of 10. That is, the trust reciprocity effect will create a horizontal asymmetry in the pattern. Therefore, we expect the line to the right of 10 to be somewhere in the shaded area. Finally, note that at the extreme observations (i.e. 0 and 20) the aggregate performance measure becomes a perfect measure of individual contributions to joint production. Therefore, we model a discontinuity at these extreme observations.

Panel B: *EXTREME* on x-axis[b]

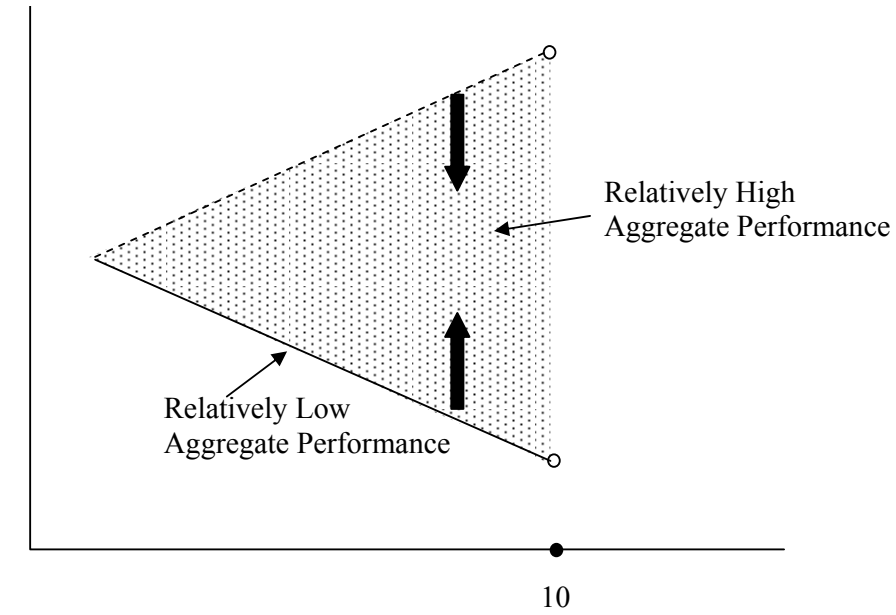

[b] Note that as EXTREME increases and aggregate performance is relatively low (i.e., aggregate performance is *decreasing* from 10 to 1), the slope of the line representing the manager's willingness to pay for information is negative, reflecting both the fairness (H2) and the trust reciprocity (H3) effects. As EXTREME increases and aggregate performance is relatively high (i.e., aggregate performance is increasing from 10 to 20), the slope of this line becomes less steep, reflecting the positive effect of trust reciprocity (H3) and the offsetting negative effect of fairness (H2). Depending on the relative sizes of these effects (which we do not predict), the slope may remain negative or become positive. Importantly, if the slope becomes positive (indicating a stronger effect for fairness than for trust reciprocity), we would expect this positive slope to be less steep than the negative slope when aggregate performance is relatively low. That is, the trust reciprocity effect will create a vertical asymmetry in the pattern. Therefore, for relatively high aggregate performance, we expect the line to be somewhere in the shaded area. Finally, note that at the extreme observation of 10, the aggregate performance measure becomes a perfect measure of individual contributions to joint production. Therefore, we model a discontinuity at this extreme observation.

FIGURE 2
Descriptive results: Mean Price Offer as a Function of the Aggregate Performance Measure

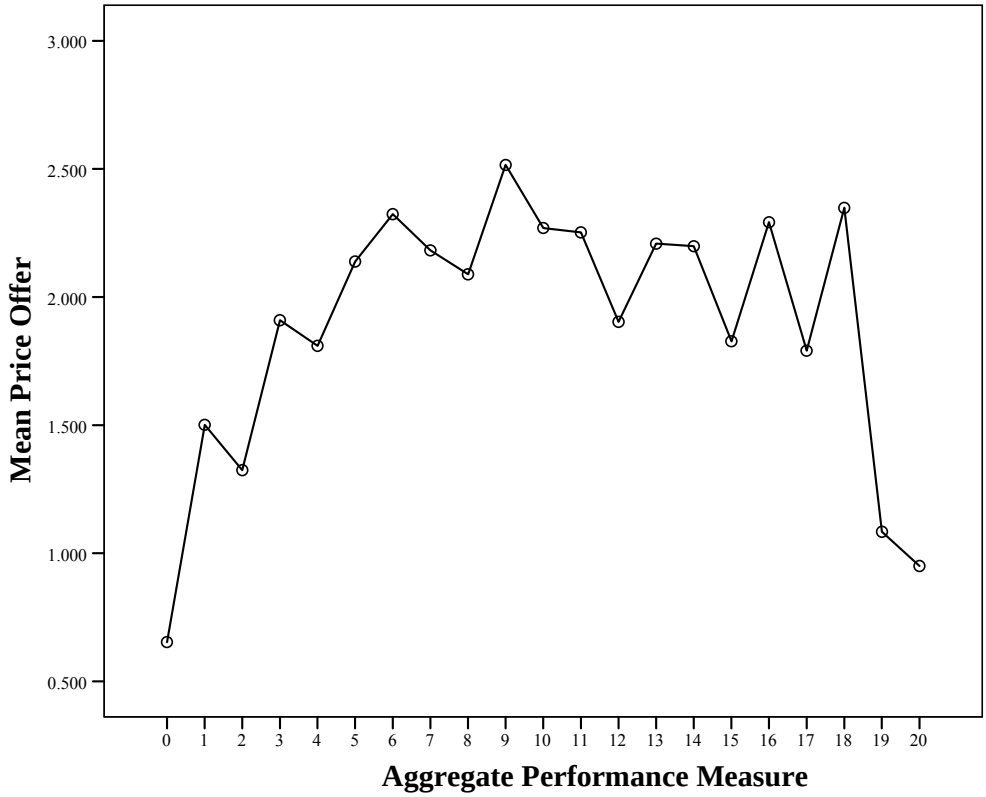

FIGURE 3
GEE Results: Willingness to Incur Cost as a Function of Aggregate Performance Measure [a]

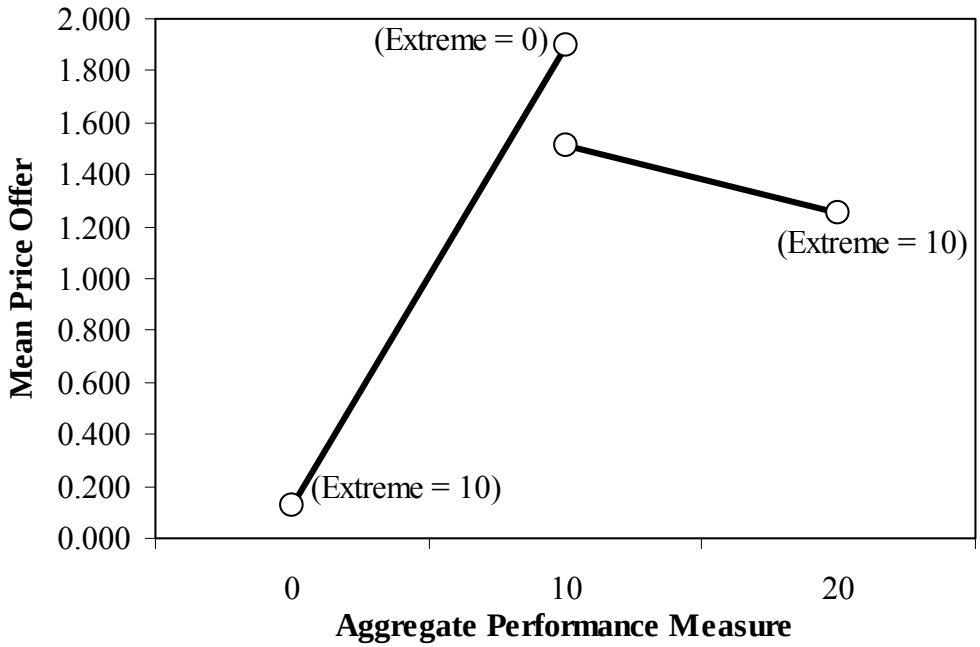

[a] This graph presents the results of the regression presented in Table 7. When aggregate performance equals 10, the variable *EXTREME* equals 0. Therefore, the regression intercept equals the price expected as aggregate performance approaches 10 from the left. This intercept plus the regression coefficient on the dummy variable *HIGHPERF* equals the price expected as aggregate performance approaches 10 from the right. While these two points are separated in the figure, it should be noted that the difference (i.e., the coefficient on *HIGHPERF*) is not significant at conventional levels ($p = 0.118$). The negative of the coefficient on *EXTREME* equals the slope of the regression line to the left of 10, whereas the sum of the coefficient on EXTREME and the interaction coefficient equals the slope of the regression line to the right of 10.

FIGURE 4
GEE Results: HIGHFAIR Managers' Willingness to Incur Cost as a Function of Aggregate Performance Measure [a]

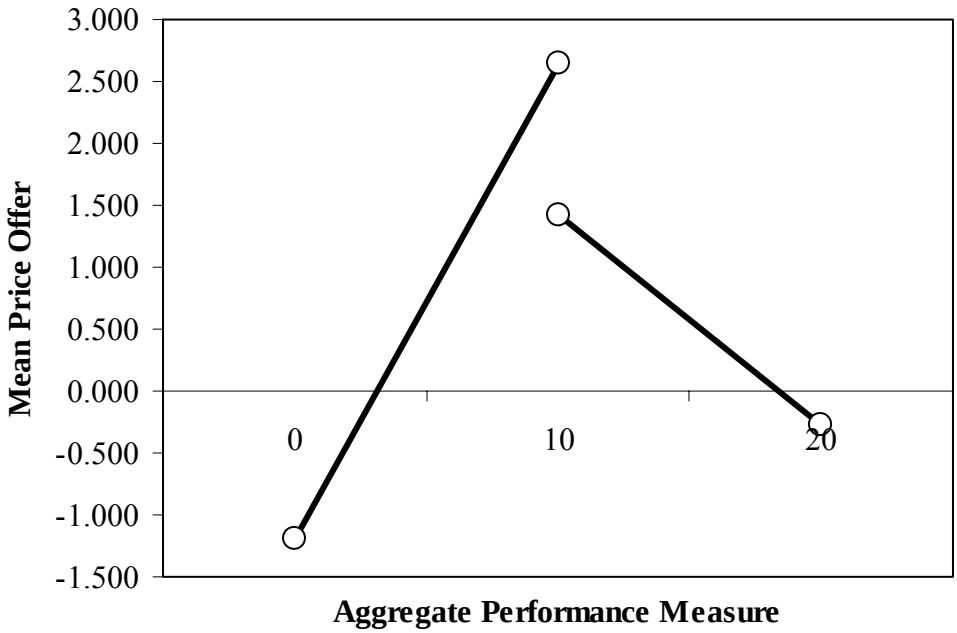

[a] This graph presents the results of the regression primary regression, including only those managers designated as HIGHFAIR. We subtract the sum of exit questionnaire items related to fairness (items 3 to 7) from the sum of those related to trust reciprocity (items 8 to 11). Those managers who score in the bottom 25% of this difference score are labeled HIGHFAIR managers, and those who score in the top 25% are labeled HIGHTRUST managers.

FIGURE 5
GEE Results: HIGHTRUST Managers' Willingness to Incur Cost as a Function of Aggregate Performance Measure [a]

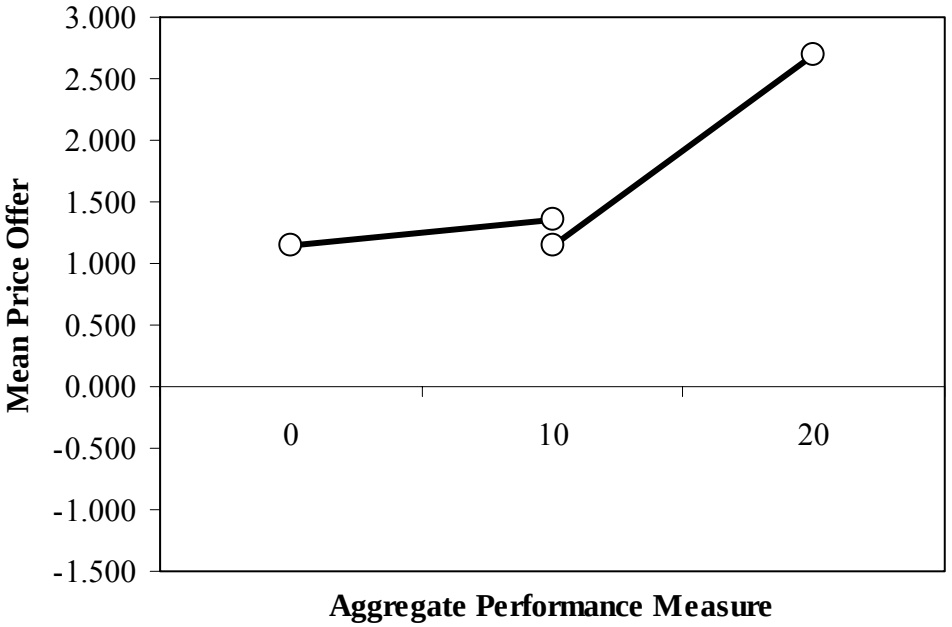

[a] This graph presents the results of the regression primary regression, including only those managers designated as HIGHTRUST. We subtract the sum of exit questionnaire items related to fairness (items 3 to 7) from the sum of those related to trust reciprocity (items 8 to 11). Those managers who score in the bottom 25% of this difference score are labeled HIGHFAIR managers, and those who score in the top 25% are labeled HIGHTRUST managers.