ACADEMIC PRESS

Games and Economic Behavior 40 (2002) 203–231

GAMES and Economic Behavior

www.academicpress.com

Altruism, equity, and reciprocity in a gift-exchange experiment: an encompassing approach

Gary Charness [a,*] and Ernan Haruvy [b]

[a] *Department of Economics, University of California, Santa Barbara, CA 93106, USA*
[b] *Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel*

Received 13 September 1999

Abstract

Considerable experimental evidence suggests that non-pecuniary motives must be addressed when modeling behavior in economic contexts. Recent theories of non-pecuniary motives can be classified as altruism-, equity-, or reciprocity-based. We outline the qualitative differences in prediction these alternative explanations yield in a gift-exchange game. We estimate and compare leading approaches in these categories, using experimental data. We then offer a flexible approach that nests the above three approaches, thereby allowing for nested hypothesis testing and for determining the relative strength of each of the competing theories. In addition, the encompassing approach provides a functional form for utility in different settings without the restrictive nature of the approaches nested within it. Using this flexible form for nested tests, we find that intentional reciprocity, distributive concerns, and altruistic considerations all play a significant role in players' decisions. © 2002 Elsevier Science (USA). All rights reserved.

JEL classification: A13; C50; C70; C91; D21; J30

Keywords: Altruism; Experiment; Equity; Fairness; Inequality aversion; Reciprocity

* Corresponding author.
E-mail addresses: charness@econ.ucsb.edu (G. Charness), eharuvy@hbs.edu (E. Haruvy).

PII: S0899-8256(02)00006-4

1. Introduction

A broad set of experimental results indicates that people frequently choose actions that do not maximize their monetary payoffs. People reject positive offers in the ultimatum game, make positive allocations to anonymous parties in the dictator game, and make voluntary public contributions that reduce their own material reward.[1] There is field evidence as well, as seen in anonymous charitable contributions. It is now generally accepted that humans often have social or psychological considerations which may lead them to sacrifice monetary payoffs in the course of maximizing their utility. This can have substantial economic consequences even for those who do, in fact, maximize only their own expected monetary payoff.

What are the underlying motivations in these cases? In the past few years there have been a number of theories attempting to explain this behavior. One explanation is simple altruism, where people care not only about their own material well-being but also about the material well-being of others. However, theories of unconditional altruism may be overly simplistic, as an agent is assumed to assign a constant weight to the welfare of every other individual. Another approach expands the altruism principle by incorporating distributive concerns. Under such models, also known as equity (or fairness) models, one's regard for another person's monetary well-being depends on the other person's monetary payoff relative to one's own. The seminal models in this category are ones of Bolton and Ockenfels (2000) (henceforth, BO) and Fehr and Schmidt (1999) (henceforth, FS). Finally, reciprocity theories expand the principle of altruism in a different direction, asserting that regard for someone else's payoff depends on how 'kind' the other is perceived to be. The principle of reciprocal altruism emphasizes the relationship between one's action and one's beliefs about the intentions of the other agent(s). A leading model in this category is one of Rabin (1993), while Levine (1998) offers an alternative formulation.

The issue of which approach, if any, best explains experimental data remains open. Each has been shown to effectively explain outcomes in only some proper subset of the various experimental settings. However, as stated by BO (p. 166), "if no connections can be found, we are left with a set of disjoint behavioral charts, each valid on a limited domain." Some have argued that one must combine the insights of altruism, distributive concerns, and reciprocity to obtain a sufficiently descriptive model. In this paper, we estimate representative models from these general approaches and suggest a simple encompassing approach to estimation which nests altruism, distributive concerns, and reciprocity. Our encompassing approach allows for some nested comparisons and testing of leading models and offers a flexible alternative that can potentially overcome some of their limitations.

[1] Roth (1995), Güth and Tietz (1990), and Ledyard (1995) offer some surveys of this literature.

2. Literature

In the ultimatum game, one person in a mutually anonymous pair proposes a division of a sum of money to the other person, who chooses to either accept the proposal or reject it. A rejection means both people receive no money. As rejections are not uncommon (particularly with rather uneven proposals), it seems that people are sometimes willing to deliberately sacrifice money. Using a sequential version of the ultimatum game, Ochs and Roth (1989) demonstrated that many players would reject a proposed positive sum of money and subsequently make counter-proposals which, if accepted, would give them less money than would the original proposal, but would still give them more money than the other player.[2]

These disadvantageous counter-offers were explained by Bolton (1991). The main feature of this model is that an individual's utility is composed of both her material payoffs and some disutility when her monetary reward is smaller than that of the other player. However, people do not mind at all if the disparity in material payoffs is in their favor. While this model explains ultimatum game behavior, it does not address results such as positive dictator game allocations and the costly effort provision found in gift-exchange experiments (e.g., Fehr et al., 1993).

BO address the asymmetry of the Bolton (1991) model by asserting that people suffer disutility whenever payoffs are unequal; it is assumed that this disutility is symmetric with respect to which player receives the greater material payoff. However, reciprocity and intention are excluded from this theory, as the authors argue that there is little laboratory evidence that these have significant influence. FS allow for an asymmetric effect of relative payoffs with a *self-centered inequality aversion*. Both models have a moderate degree of success in explaining behavior in bilateral bargaining, public goods environments, and markets.

Notwithstanding their success in some settings, distributional motives alone fail to explain results in other settings. For example, Kahneman et al. (1986) find that three-quarters of all participants are willing to sacrifice $1 and thereby punish a person who chose a selfish allocation in a previous dictator game and reward a person who chose a generous one.[3] Blount (1995) elicits the minimum acceptable offer in a strategy-method version of the ultimatum game. A comparison between treatments where the offer is generated by either a random mechanism or a self-interested party shows a clear difference in the willingness to accept lopsided offers. Offerman (1998) varies the attribution for an outcome and observes considerable negative reciprocity, but limited positive reciprocity.

[2] There was a substantial discount factor between periods.

[3] The choice was between ($6, $0, $6) and ($5, $5, $0).

Kagel et al. (1996) vary the exchange rates for payoff chips and the information provided about these exchange rates. They find that ultimatum rejection rates depend on responder beliefs about proposer knowledge of the exchange rates, as intentional (informed) unequal proposals were rejected at substantially higher rates than unintentional unequal proposals. Brandts and Charness (1999) test for punishment and reward in a cheap-talk game and find that intention is a critical issue, finding substantial negative reciprocity and significant, but limited, positive reciprocity. Andreoni et al. (1999) find that “fairness is a function of more than just the final allocations of subjects, but depends on the actions that were not chosen as well as those that are.”

The psychology literature was first to include reciprocity in a model of human behavior. Heider (1958) introduced the idea that causal inference, where one takes into account another person’s motives and situational constraints, is an important cognitive process for perceiving social contexts. Individuals have a need to infer causes and to attempt to assign responsibility for outcomes. When volition is absent, feelings of revenge and gratitude dissipate or vanish. Experimental studies include Greenberg and Frisch (1972), who find that deliberate help leads to more reciprocity than does accidental help, and Goranson and Berkowitz (1966), who confirm a differential sequential response conditioned on prior help provided by another agent.[4]

Rabin (1993) suggests that intentional reciprocity, both positive and negative, can play a role. However, the model presented in the main text of his paper does not explain the ‘reciprocity-free’ component of distributive concerns, which seems to be present in many experimental results. For example, this model does not explain positive allocations in the dictator game. Levine (1998) combines both general altruistic or spiteful tendencies with a personal component which accommodates a form of reciprocity. However, this model does not permit any influence for comparative payoffs and the range of the non-pecuniary term is rather limited.

3. Experiment

We examine data from a gift-exchange game (Charness, 1996). A total of 122 subjects participated in this experiment; 61 each in the ‘employer’ and ‘employee’ roles. Average earnings (including a $5 show-up fee) were between $16 and $17 in a session of less than 2 hours. All subjects were students at the University of California at Berkeley. The design was quite similar to the BGE/TC treatment by Fehr et al. (1998); the only change was that employees could not reject the wage.[5]

[4] See also Thibaut and Riecken (1955), Kelley and Stahelski (1970), Kahn and Tice (1973), and Thomas and Batson (1981).

[5] The complete instructions can be obtained from http://www.econ.upf.es/cgi-bin/onepaper?283.

There were, in general, ten employers and ten employees in each session, where a session consisted of ten periods. Employers and employees initially all met in one large room. Pairings were anonymous and it was common knowledge that employees and employers were not rematched.[6]

An employee was given a wage, which had been assigned by either an employer or an external process. The process by which the wage was determined was fully known to the employee. No employee was in more than one of the treatments, since each session was a different treatment. Once assigned a wage, each employee was asked to record her effort choice (between 0.1 and 1.0, inclusive) on a record sheet. The monitor then gave this sheet to the corresponding employer in the other room, so that the employer had physical evidence of the employee's choice. This procedure was common knowledge.

The combination of wage and effort determined outcomes and monetary payoffs for each pair of subjects in a period. Each employer was given an endowment of 120 "income coupons" in each period. The monetary payoff functions were given by

$$\Pi_{\mathrm{F}} = (120 - w) \cdot e, \tag{1}$$

and

$$\Pi_{\mathrm{E}} = w - c(e) - 20, \tag{2}$$

where Π_{F} denotes the payoff to the employer (firm), Π_{E} denotes the payoff to the employee, e denotes the employee's effort, w is the wage, and $c(e)$ is the cost of effort, a non-linear function increasing in e. The exact cost function facing an employee is as follows:

Effort	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Cost	0	1	2	4	6	8	10	12	15	18

The innovation in this experiment was that the source of the wage differed across sessions and was known to be generated by either a self-interested employer, a draw from a bingo cage, or the experimenter (in advance). Only one wage-determination mechanism was used in each session.

Note that the unique Nash equilibrium (which can be easily deduced by backwards induction) of $w = 20$ and $e = 0.1$ is socially inferior as there are many outcomes which would make both players better off. Indeed, employees and employers consistently deviate from their Nash equilibrium choices, with employers choosing on average wages in excess of 50 and employees choosing on average effort levels greater than 0.3.

[6] In some cases, there were less than 20 subjects in a session, so that there was inevitably some rematching. However, these were still anonymous and subjects were assured that no pairing would be repeated in two consecutive periods. An analysis of the data (Charness, 1996) shows that there were no apparent reputation effects.

4. Theory

This section is divided into five sections. Each of the first three sections presents a simple representative model corresponding to a general approach. The fourth section outlines the qualitative differences in prediction between the three approaches. The last section presents an approach that nests the three mainstream approaches, thereby providing a useful tool for comparison of behavioral motives.

It is common in experimental design of this nature to estimate only the second mover's utility function. Since estimating a utility function for the first mover, the employer, involves specifying expectations on the second mover's actions, it is easier (if we are willing to relax the assumption of equilibrium beliefs) to estimate only the employee's utility function based on her choices, taking the employer's action and the employer's perception of the employee's intentions as given. Hence, all utility functions in the sections that follow will be defined for the employee only.

4.1. Altruism

The basic altruism model in this setting would postulate a utility function for the employee of the form

$$U_{\mathrm{E}}(\Pi) = \Pi_{\mathrm{E}} + \alpha \Pi_{\mathrm{F}}, \tag{3}$$

where α, the altruism parameter, is constant over the payoff space in this formulation. Altruists are considered to have $\alpha > 0$, where α equals 0 for pure money-maximizers. It is also generally accepted that $\alpha \leqslant 1$. In other words, another person's monetary payoff is not more rewarding than one's own. While each individual has a distinct value for α, this value is independent of previous actions or history for other agents. Costa-Gomes and Zauner (2001) called the above formulation "social utility" and estimated it on four-country ultimatum game data, finding that it had good explanatory power.

4.2. Fairness and distributive justice

4.2.1. The egalitarian notion of fairness

The egalitarian notion of fairness implies that a fair distribution assigns everyone an equal share of resources. Two simple models corresponding to this notion are BO and FS. Although there are some qualitative differences between the BO and FS models, they generally make similar behavioral predictions, particularly in the case of two agents. We use the FS model for purposes of comparison. Applying the two-agent version of this model to the gift-exchange setting, we have

$$U_{\mathrm{E}}(\Pi) = \Pi_{\mathrm{E}} - \alpha\big[\max(\Pi_{\mathrm{F}} - \Pi_{\mathrm{E}}, 0)\big] - \beta\big[\max(\Pi_{\mathrm{E}} - \Pi_{\mathrm{F}}, 0)\big], \tag{4}$$

where it is assumed that $\beta \leqslant \alpha$ and $0 \leqslant \beta < 1$. This means that, in addition to being concerned with one's own pecuniary reward, one also cares about equity. Furthermore, assuming $\beta > 0$, one's disutility is increasing in the distance between payoffs. This parsimonious model captures the concept that one cares about one's payoff not only in absolute terms but also in comparison to others. However, a potential drawback (due to the non-negativity restriction on the model's parameters) is that this model does not nest simple altruism, where one receives some positive utility from the level of another agent's material payoff, independently of one's own.

Another drawback is that often α and β cannot both be estimated. Note that in ultimatum games only α can be estimated.[7] In the gift-exchange context, only β can be meaningfully estimated, as the only way for an employee to sacrifice monetary payoffs when assigned an unfavorable wage is to increase her effort, thereby *rewarding* the employer for his unkind behavior!

A third concern involves the linearity of one's utility in the payoff difference. It may not be reasonable to expect a sharp discontinuity, at the reference point, in the weight one assigns to the other's payoff. Furthermore, it may be unreasonable to expect one's weight in a close neighborhood of the reference point to be the same as when the outcome reflects substantial inequity. A quadratic function in the payoff difference might seem more plausible. However, the first order condition with such a representation makes the problem rather intractable. Moreover, coefficients lose the interpretation of weight on inequality.

4.2.2. Other notions of fairness

Two common alternative concepts of distributive justice are the utilitarian and Rawlsian social welfare functions. The utilitarian view asserts that people should maximize some function which is monotonically increasing in the well-being of each member of society. The simplest utilitarian model maximizes the sum of payoffs (Eq. (3) with $\alpha = 1$). On the other hand, the Rawlsian approach involves maximizing the payoff of the worst-off agent.

An experiment by Charness and Rabin (2001) sheds some light on these differing notions of equity. In that experiment, a third player (henceforth, player C) was to choose an outcome for two anonymous players. In one treatment, C was asked to choose between (1200, 0) and (400, 400) for players (A, B); here 18 out of 22 (82%) subjects picked (400, 400). Yet, in the second treatment, when C chose between (750, 375) and (400, 400), only 11 out of 24 (46%) selected (400, 400). The difference was significant at the 1% level. In the first case, the egalitarian and Rawlsian notions appeared to have better explanatory

[7] In order to be able to estimate one's weight on inequality, some inequality-reducing action must be observed that goes counter to money-maximizing behavior. One rarely observes a responder rejecting an offer (the only way to sacrifice monetary payoffs) which would favor her in relative material payoffs. Such a rejection would amount to punishing the proposer for his apparent generosity!

power relative to the simple utilitarian model; however, this was reversed in the second case.

Thus, it appears that equity is not limited to considerations of equality. Although agents' relative shares may be important, so is the total amount to be distributed. This suggests that inequality aversion is not the only proxy for fairness. As we find in Section 6, the FS model does not fare well with the gift-exchange results. Nonetheless, distributive and altruistic concerns are the only plausible explanations for non-minimal effort provision in the non-volitional treatments, since the employer's lack of choice means no motives of kindness or unkindness can be attributed.[8]

4.3. Intentional reciprocity

A critical question concerns the effects of the perceived intentions of others on behavior. Suppose an agent cares about money and comparative payoffs, but also has feelings (*regard*) about another agent's absolute payoffs. If this regard for another person depends on one's perceptions of that person's prior actions, we would require a model of reciprocity, or *reciprocal altruism*, as exemplified by Rabin (1993).[9] The key to this model is that one's utility is affected by one's beliefs about another agent's motivation for making a particular choice. A potential drawback is that reciprocity models do not generally nest or explicitly incorporate concerns for relative payoffs or altruistic behavior.

We present the framework of Rabin (1993).[10] The employee's kindness to the employer as a function of effort is given by

$$f_{\mathrm{E}}(e|w) = \frac{\Pi_{\mathrm{F}}(e|w) - \Pi_{\mathrm{F}}^{e}(e|w)}{\Pi_{\mathrm{F}}^{\max}(e|w) - \Pi_{\mathrm{F}}^{\min}(e|w)}, \tag{5}$$

where $\Pi_{\mathrm{F}}^{e}(e|w) = (\Pi_{\mathrm{F}}^{\max}(e|w) + \Pi_{\mathrm{F}}^{\min}(e|w))/2$ is the suggested reference point.[11] Given that all points are Pareto-efficient, $\Pi_{\mathrm{F}}^{\max}(e|w)$ denotes the maximum that the employer could make given the wage w. Similarly, $\Pi_{\mathrm{F}}^{\min}(e|w)$ denotes the minimum that the employer could make given the wage w.

To gauge the employer's kindness (as perceived by the employee, a critical issue here), we would need to know the employee's belief about the employer's

[8] Recall that altruism is a subset of the utilitarian notion of fairness.

[9] While the normal-form Rabin model does not apply formally to the Akerlof and Yellen gift-exchange context, Dufwenberg and Kirchsteiger (1998) extend this model to sequential games and get qualitatively similar predictions in almost all cases.

[10] Rabin's framework is an equilibrium framework. It is crucial to emphasize that the analysis at hand borrows only the functional form for the player utility from Rabin (1993) and not the equilibrium framework, which is not straightforward to apply here.

[11] This reference point may be inappropriate in an equilibrium model. However, Π_{F}^{e} drops out in the first order condition of maximization with respect to e.

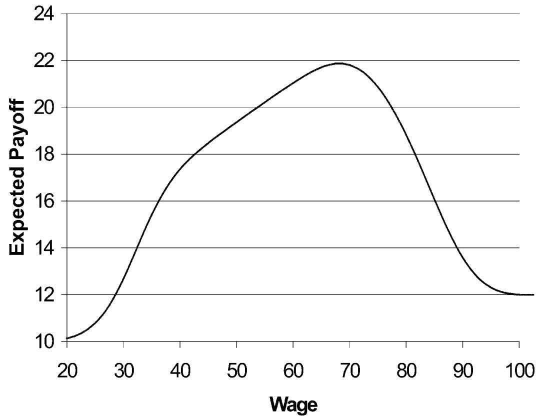

Fig. 1. Employer's expected payoff on wage, done by kernel regression.

belief about the relationship of effort to the wage assigned. For example, if the employee believed that the employer expected the same effort regardless of the wage, any wage above 20 would be considered "kind." However, if the employee believes that the employer is merely offering a higher wage in an attempt to maximize her own expected earnings, one might consider a kind wage to be one that exceeds the profit-maximizing level.

As we are unable to observe an employee's beliefs, we shall assume that he knows that the employer knows that this profit-maximizing wage is 70, which is *ex post* seen to have actually yielded the employer the highest expected profits, as is shown in Fig. 1.

Incidentally, a wage of 70 is the wage that would yield the middle point in the opponent's payoff given the opponent's action, the reference point suggested by Rabin (1993). Hence, the employer's kindness is now gauged by

$$f_{\mathrm{F}}(e|w) = \frac{w - w^{e}}{w^{\max} - w^{\min}}, \tag{6}$$

where $w^{e} = 70$. Thus, the values of f_{F} and f_{E} must lie in the interval $[-\frac{1}{2}, \frac{1}{2}]$. These kindness functions can now be used to specify the employee's preferences

$$U_{\mathrm{E}} = \Pi_{\mathrm{E}} + f_{\mathrm{F}} f_{\mathrm{E}}. \tag{7}$$

The central behavioral feature of this reciprocity-based utility specification [hereafter Reciprocity I] reflects the main principle in models of reciprocity: If the employee believes that the employer is treating him poorly ($f_{\mathrm{F}} < 0$) then he will reciprocate with a negative value for f_{E} and would attempt to reduce the employer's payoff. On the other hand, if the employee believes that the employer is treating him favorably ($f_{\mathrm{F}} > 0$), then he prefers $f_{\mathrm{E}} > 0$ and will sacrifice money to benefit the employer, leading to a non-minimal choice of effort.

A variation on this functional form, presented in the appendix of Rabin (1993), nests altruism as a special case. This formulation, henceforth Reciprocity II, is

$$U_E = \Pi_E + \left[\alpha + (1-\alpha) f_F\right] f_E. \tag{8}$$

In this utility function, when $\alpha > 0$ the employee will wish to be kind to the employer, even when the latter is "neutral" to the employee. If α is close to 1, then pure altruism dominates behavior. Because the kindness functions are bounded above and below, the behavior in this model is sensitive to the scale of material payoffs.

Rabin's specification is qualitative in nature and does not provide the relative power of non-pecuniary versus material interests. However, Rabin acknowledges that the non-pecuniary term in the utility function must be rescaled to correspond to the payoff scale. Here the scaling factor should be proportional to the range of payoffs for the employer, so the applicable factor is the payoff range for the employer, which is $0.9 \times (120 - w)$. Besides rescaling the non-pecuniary term to the employer's payoff scale, this scaling factor has the desirable feature that when the employer is neutral, the model reverts to the simple altruism model of Eq. (3) and α is the same as the altruism parameter therein. Given the specification of f_E and f_F and the rescaling by $0.9 \times (120 - w)$, we get

$$U_E = \Pi_E + \left[\alpha + (1-\alpha)\left(\frac{w-70}{100}\right)\right](e - 0.55)(120 - w). \tag{9}$$

4.4. *Qualitative differences in prediction*

It is reasonable to ask how the three theories discussed above would result in substantial differences in predicted behavior in the setting at hand.

We begin by comparing altruism to distributive justice as defined by the FS formulation. Whereas both predict positive effort levels, the relationship between effort and wage is radically different under the two theories in the setting at hand. In the gift-exchange setting, the incremental benefit to the employer from additional effort is decreasing in wage. For example, when the wage is 50, an increase in employee effort by an increment of 0.1 will increase the employer's payoff by 7. When the wage is 80, the same increase in effort by the same increment will result in only 4 additional units for the employer. Given the fixed weight in the employee's utility on the employer's payoff under the altruism explanation, effort is non-increasing in wages (and is decreasing in wage at any point in the wage range where effort is greater than the minimum). In other words, *under altruism, the relationship between effort and wage is negative*.

The FS equity notion would predict the opposite effort-wage relationship. For low wages, the employee would not attempt to help the employer before he receives at least as much as the employer (at any wage below 30, the employee would not choose costly effort, since the employer receives more than the employer even with minimum effort). Since lower wages would require the "equality-generating effort" to be lower, *under equity, the relationship between effort and wage is positive*.

Both the equity and reciprocity models can potentially explain the positive effort-wage relationship. They both specify (explicitly or implicitly) a reference point in the employee's utility function. The difference lies in the nature of the reference point. The FS model specifies a reference point with equal payoffs regardless of the wage chosen. On the other hand, the reciprocity specification for a reference point has the employer's relative payoff increasing in wage, since the employer's appropriate relative share changes in the employer's favor with higher wages. In functional terms, this distinction results in a quadratic relationship between effort and wage in the reciprocity model, versus a linear relationship in the FS equity model.

Perhaps an equally important distinction is where the base reference point (in the reciprocity model, the implicit reference with no kindness) should be. Whereas FS equity specifies it at the equal share point, reciprocity suggests an implicit base reference point more dependent on the action space and corresponding payoff range than on the distribution of payoffs.[12] The reciprocity approach can therefore be more flexible in settings with inherent inequality in the initial endowments, as in the case at hand.

4.5. An encompassing approach

While all of these models have a considerable degree of success in explaining deviations from monetary payoff maximization, each fails to plausibly account for a non-trivial proportion of observed behavior. One relatively parsimonious combination of these models is

$$U_{\mathrm{E}}(x) = \Pi_{\mathrm{E}} + R(w)\Pi_{\mathrm{F}}, \tag{10}$$

where $R(w)$ represents the *regard* of the employee for the employer as a function of the employer's action.

Definition. *Regard* is defined as the weight in one's utility function on the monetary payoff of the other. Regard is formulated as a function of the state of the world (initial endowments) and the other's (observed or expected) action.

Claim. *When the regard is formulated as* $R(w) = \delta D + a + bw + cw^2 + dw^3 + ew^4$*, where* (δ, a, b, c, d, e) *are parameters to be estimated and* D *is a dummy variable taking the value of* 1 *when* $\Pi_{\mathrm{E}} > \Pi_{\mathrm{F}}$ *and* 0 *otherwise, pure altruism, Fehr–Schmidt equity, and the reciprocity specification of Eq.* (9) *are all special cases of this formulation.*

[12] Of course, the main-text Rabin (1993) model is distribution-free and predicts only minimum effort in the absence of volition wages.

(I) Imposing the restriction that $\delta = b = c = d = e = 0$, we get the pure altruism case of Eq. (3) with $\alpha = a$.
(II) Imposing the restriction that $b = c = d = e = 0$, we get the Fehr–Schmidt formulation of Eq. (4) with $\alpha = -a/(1 + a)$ and $\beta = (a + \delta)/(1 + a + \delta)$.
(III) Imposing the restrictions that $b = (1 - a)/170$ and $\delta = c = d = 0$, we get the Rabin formulation of Eq. (9) with $\alpha = 1 - 100b$.

Therefore, we can use a simple χ^2 test to compute a significance level for each approach relative to the encompassing model and compare the approaches' strengths relative to the encompassing model as well as their strengths relative to each other when nestedness is possible. In addition, our approach by itself provides a flexible alternative to assess a functional form in different settings without the restrictive nature of the approaches nested within it. For example, when there is no intention (as is the case of the non-volition treatments), the encompassing model is better suited to examine distributive concerns since the linearity restriction on the payoff difference is eliminated. Theoretical justification aside, $R(w)$ can be thought of as a flexible fourth-degree polynomial approximation of an unknown function in the wage.

An important difference from the conventional (Rabin) reciprocity model is that $R(w)$ is not necessarily 0 in the default case, when the employer has had no choice to make.

5. Preliminary tests and estimation procedure

We begin by examining whether we can pool observations over time periods. While in principle the matching protocol prevents reputation effects from forming, players may be nevertheless modifying their behavior over time. Looking at the average effort, wage, and effort/wage ratio for each period (Appendix B), there are no apparent changes in behavior. Simple joint hypothesis F-tests (Appendix B) confirm this assertion. This is consistent with similar findings by Fehr and co-authors in gift-exchange experiments.

Next we examine whether the two non-volition treatments (bingo cage draws and experimenter-determined wage) can be pooled with each other and whether they can be pooled with the volition (employer-determined wages) case. If we were to find no significant difference between behavior in volition and non-volition sessions, we would be able to eliminate intentional reciprocity as a factor.[13] For that purpose, we adopt simple flexible parametric testing procedures. Each treatment's effort levels can be estimated as a function of effort on wage. A flexible functional form is the polynomial function. A fifth-degree

[13] Charness (1996) finds a significant difference using a linear Tobit model.

polynomial was used (higher-order polynomials did not significantly improve likelihood), of the form: effort $= \sum_{i=0}^{5} c_i \cdot \text{wage}^i$. This functional form was estimated for each of the treatments using the Tobit procedure (effort is left-censored at 0.1, with numerous observations of $e = 0.1$). Next, the two non-volition treatments were pooled and the parameter estimates from the pooled regression were imposed on each of the three sessions. The statistic for the likelihood ratio test is $\chi^2 = -2\ (\log L_\mathrm{r} - \log L_\mathrm{u})$, where L_r is the restricted likelihood and L_u is the unrestricted one. Under the null, this statistic is distributed with degrees of freedom equal to the number of restrictions (in our case, 6). The statistics for the three treatments are

Bingo-cage	$-2*[(-99.1221)-(-99.1055)] = 0.0332$
Experimenter-determined	$-2*[(-78.6198)-(-78.2930)] = 0.6536$
Employer-determined	$-2*[(-117.334)-(-87.0608)] = 60.547$

Critical value at the 5% level of significance is 12.592.

Hence, we cannot reject pooling the two non-volition treatments at any reasonable level of significance, but we can easily reject pooling across volition conditions at all reasonable levels of significance. Proceeding in this manner, we obtain Fig. 2 (Predicted effort levels below 0.1 are not in the range of efforts available to subjects and are due to the probit portion of the tobit regression.)

Comparing across treatment groups, the effort levels look roughly the same in the center of the wage range. However, we see substantial differences at the extremes of the wage range. This would seem to offer an opportunity to separate equity and reciprocity considerations. The area in between the curves reflects intentional reciprocity since intention is the only difference between the two treatments. Yet, one cannot readily draw conclusions since no structural form for utility has been specified. Hence, we cannot separate non-pecuniary concerns

Fig. 2. A Tobit regression of effort on wage, using a fifth-degree polynomial.

from pecuniary ones, much less equity from reciprocity. One must estimate a functional form for utility.

Still, there are a few observations worth noting in Fig. 2:

(1) Non-volitional effort is rising at an increasing rate at wages above 85. This is counterintuitive since effort is costly at high effort levels whereas the benefit to the employer from higher levels of effort is small at such wages. We may attribute this to increasing importance on the opponent's pay due to non-linear distributive concerns. Alternatively, this phenomenon can be explained by an equity notion defined by effort rather than by pay levels.
(2) At low wages (less than 35), effort predictions are below 0.1. Since we do not have observations on effort levels below 0.1, these predictions arise from the Probit portion of the Tobit procedure.

Let ν denote the vector of parameters to be estimated for a given model. Then for a given model the functional form for utility can be represented by $U_{\mathrm{E}}(w, e; \nu)$. The first step is to differentiate $U_{\mathrm{E}}(w, e; \nu)$ with respect to e. To do so we must first define a smooth approximation for $c(e)$. The quadratic form is

$$c(e) = 9.9(e - 0.1) + 11.1(e - 0.1)^2. \tag{11}$$

The first-order condition for maximization would be

$$\frac{\mathrm{d}U_{\mathrm{E}}(w, e; \nu)}{\mathrm{d}e} = 0. \tag{12}$$

Solving (12) for e, we get the optimal effort e as a function of w and ν

$$e = g(w; \nu). \tag{13}$$

We now define the likelihood function to be maximized with respect to ν as

$$L = \prod_{i=1}^{N} \prod_{t=1}^{T} \phi\bigl(e_{it} - g(w_{it}; \nu)\bigr)^{D_{it}} \Phi\bigl(0.1 - g(w_{it}; \nu)\bigr)^{1-D_{it}}, \tag{14}$$

where D_{it} (i indexes individuals and t indexes time) is a dummy that takes the value of 1 if effort exceeds 0.1 and 0 otherwise.

6. Estimation results

Detailed estimation results are given in Tables 3–13. (Tables 3–8 provide parameter estimates for the various models under the volitional treatment, Tables 9–13—under the non-volitional treatments.) Below we present a summary for each model discussed above. The parameter σ is obtained for each model and denotes the standard deviation of the error term. We report the standard deviation as a measure of goodness of fit. We also compare the log-likelihood over models.

Table 1
The employer volition sessions (220 observations)

Model	σ	Log-likelihood
Pure pecuniary	0.779	−187.1
Simple altruism	0.612	−174.3
Fehr–Schmidt (w/only β)	0.600	−165.5
Fehr–Schmidt (w/both α & β)	0.600	−165.5
Reciprocity I	1.128	−211.7
Reciprocity II	0.353	−99.9
Unrestricted encompassing	0.334	−90.2

Note that all models are nested within the encompassing (flexible polynomial) approach.

Given the above values for σ and the log-likelihood for each model, one can make some nested comparisons. We observe that:

(1) Pure pecuniary considerations can be rejected in favor of pure altruism, indeed in favor of any one of the models estimated, except for Reciprocity I.
(2) Though altruism is successful relative to a model of pure monetary considerations, it can be rejected in favor of (an unrestricted[14] two-parameter) FS formulation.
(3) The two-parameter formulation of FS does not significantly outperform a one-parameter version here. This is expected given the inability by players to sacrifice payoff for equality when they are already monetarily disadvantaged relative to others.
(4) Altruism can also be rejected in favor of the Reciprocity II version, showing promise for Rabin's utility specification as an explanation. As we will shortly demonstrate, this success may be attributed to other factors. Nonetheless, altruism alone cannot account for the observed behavior.
(5) The Reciprocity I model can easily be rejected in favor of the Reciprocity II formulation, demonstrating the relative importance of altruism in the determination of non-pecuniary behavior.
(6) All four models of non-pecuniary motives can be rejected in favor of the flexible polynomial approach. This implies that Reciprocity I, Reciprocity II, and FS are over-restrictive. However, a simple likelihood comparison reveals that Reciprocity II is substantially less so.

While restrictive forms result in a lower fit, they are often desirable due to parsimony and theoretical appeal. If we insist on a restrictive form, Reciprocity II

[14] By unrestricted, we mean that we relax the FS restrictions that $\beta \leqslant \alpha$ in Eq. (4).

Table 2
No employer volition (390 observations)

Model	σ	Log-likelihood
Pure pecuniary	0.782	−353.9
Simple altruism	0.516	−292.1
Fehr–Schmidt	0.603	−314.7
Reciprocity I	1.180	−433.0
Reciprocity II[a]	0.343	−190.8
Unrestricted encompassing	0.323	−181.7

[a] The non-volitional representation of Reciprocity II ought to reduce to the simple altruism model, given the impossibility of observing employer kindness or unkindness. Nonetheless, we estimate this volition-dependent formulation in the non-volitional setting for the purpose of testing to what extent the "reciprocity term" in the volition formulation truly captures reciprocity, as opposed to fairness.

does best[15] in the volition case. However, one should be cautious in interpreting this finding as supporting the notions of reciprocity and altruism, as opposed to equity considerations. The extent to which the *reciprocity term* in Rabin's utility specification captures equity consideration can be assessed by estimating the same functional form of Reciprocity II for both volition and non-volition.

If we desire a more flexible approach, the polynomial is not merely significantly better in likelihood; it allows a comprehensive approach which nests all of the restricted models, allows the relaxation of restrictions, and can potentially allow separating different motives. Parameters estimated under the flexible approach show that the equality dummy parameter (δ) is not significantly different from zero, dealing a further setback to the notion of equality as a social reference point for equity.

From the estimates under the non-volition treatments, we observe that:

(1) The Fehr–Schmidt model of equity does not explain the data significantly better than altruism.
(2) All models can be rejected by nested comparison in favor of the flexible approach.
(3) In likelihood terms alone, the Reciprocity II (volitional) formulation best fits the non-volitional treatments.

Given result (3) it would appear that the Reciprocity II functional form picks up some elements of distributive concerns as opposed to reciprocity alone. Our rejection of pure equality considerations (FS) is perhaps best explained by recalling our assertion (Section 4.2.1) that equity is not the same as equality.

[15] Only FS is not nested within Reciprocity II among the restrictive models. The difference in log-likelihood between Reciprocity II and FS is so large that a simple log-likelihood comparison, using Akaike information criterion, is sufficient in this case.

Table 3
Unrestricted encompassing model[a]

Parameter	Estimate	Std. error	t-Statistic
a	−0.704	1.951	−0.361
b	0.033	0.131	0.249
c	−2.57E–04	3.16E–03	−0.081
d	−2.23E–06	3.30E–05	−0.067
e	3.12E–08	1.25E–07	0.249
δ[b]	–	–	–
σ	0.338	0.025	13.545
	Log-likelihood −91.471		

[a] δ is restricted to the positive range, in accordance with the Fehr–Schmidt model.
[b] Restricted to the positive real numbers, in accordance with FS. Not found to significantly differ from 0.

Table 4
Altruism

Parameter	Estimate	Std. error	t-Statistic
α	0.104	0.017	6.082
σ	0.612	0.055	11.147
	Log-likelihood −174.257		

Table 5
One-parameter FS formulation

Parameter	Estimate	Std. error	t-Statistic
β	0.140	0.018	7.638
σ	0.600	0.043	13.852
	Log-likelihood −165.475		

Table 6
The two-parameter Fehr–Schmidt formulation

Parameter	Estimate	Std. error	t-Statistic[a]
β	0.140	0.018	7.638
α	0.181E–04	0.029	0.629E–03
σ	0.600	0.043	13.852
	Log-likelihood −165.475		

[a] The α parameter is restricted to be positive. The t-statistic shows that α is not significant. Hereafter, only the one-parameter formulation will be presented under Fehr–Schmidt. However, the test for significance was conducted for each treatment and α is never significant.

Table 7
Reciprocity I

Parameter	Estimate	Std. error	t-Statistic
σ	1.128	0.139	8.109
	Log-likelihood −211.67		

Table 8
Reciprocity II

Parameter	Estimate	Std. error	t-Statistic
α	0.314	0.0105	30.035
σ	0.353	0.022	16.370
	Log-likelihood −99.86		

Table 9
Unrestricted encompassing model

Parameter	Estimate	Std. error	t-Statistic
a	0.223	0.293	0.762
b	−0.024	0.028	−0.869
c	9.99E–04	8.87E–04	1.127
d	−1.46E–05	1.18E–05	−1.234
e	7.77E–08	5.60E–08	1.387
δ[a]	–	–	–
σ	0.323	0.0156	20.768
	Log-likelihood −181.733		

[a] Restricted to the positive real numbers, in accordance with FS. Not found significantly different from zero.

Table 10
Altruism (not nested within the one-parameter FS formulation)

Parameter	Estimate	Std. error	t-Statistic
α	0.156	0.010	15.828
σ	0.516	0.034	15.408
	Log-likelihood −292.099		

Table 11
One-parameter FS formulation

Parameter	Estimate	Std. error	t-Statistic
β	0.143	0.013	11.021
σ	0.603	0.033	18.385
	Log-likelihood −314.696		

Table 12
Reciprocity I

Parameter	Estimate	Std. error	t-Statistic
σ	1.180	0.061	19.330
Log-likelihood -432.972			

Table 13
Reciprocity II

Parameter	Estimate	Std. error	t-Statistic
α	0.328	0.006	59.640
σ	0.343	0.016	22.101
Log-likelihood -190.810			

It appears that the Reciprocity II formulation is capturing an equity notion not captured by FS.

Consider the following idea

$$U_E = \Pi_E + \alpha \Pi_F + \beta \left| e - e^{\min} \right| \left[\Pi_E^{\max} - \Pi_E^{\min} \right]. \tag{15}$$

In this approach (henceforth the equity formulation) α is the simple altruism parameter. The parameter β is the weight on equity, where equity is measured as the distance of effort from the minimum, and least equitable, level of effort.[16] $[\Pi_E^{\max} - \Pi_E^{\min}]$ simply rescales the equity term so that it is proportional to the range of payoffs of the employee. In other words, the employee's attention to the fairness of his action does not diminish as his payoffs increase. Wage enters into the equity term only to the extent that it affects the range of payoffs. While the resulting formulation resembles Rabin's, the interpretation is somewhat different.[17]

[16] Alternatively, using Rabin's notion of equity, the equitable effort is that effort which would give the employer the payoff that is midpoint between the maximum possible and the minimum possible, taking wage as given. That formulation gives a significantly worse fit as measured by likelihood.

[17] Rabin's formulation implies that equity is considered in *monetary terms* and attention to equity is *dependent on wage*. The above formulation of equity implies that equity is considered in *action terms* and that attention to equity is *independent of wage*.

Parameter	Estimate	Std. error
Equity formulation with volition		
α	0.092	0.014
β	0.324	0.023
σ	0.340	0.026
Log-likelihood −107.0		
Equity formulation without volition		
α	0.114	0.008
β	0.301	0.018
σ	0.323	0.015
Log-likelihood −184.5		

Notice that while in the case of volition, Reciprocity II somewhat outperforms the equity formulation, as seen by the log-likelihood comparison of −99.9 (Reciprocity II) vs. −107 (equity formulation), in the non-volition case this is reversed with a log-likelihood comparison of −184.5 vs. −190.8. Note that in both the volition and non-volition cases, the standard deviation, σ, is smaller in the equity formulation: 0.340 vs. 0.353 in the volition case and 0.323 vs. 0.343 in the non-volition case.

6.1. Robustness of parameter estimates

To test the predictive power of the flexible model out-of-sample, we estimate the model's parameters on a subset of games and use these estimates to predict the behavior in another subset of games. For this purpose, the composition of the two subsets must be similar and the subset used to estimate the parameters for prediction must be large enough to have a reasonable efficiency of parameter estimates. For the predicted subset (subset II), we chose the last five players from the bingo-cage treatment (by player index) and the first five players (by player index) from the treatment where wages were experimenter-determined. We refer to the remaining subset of 29 players in the non-volition treatments as subset I.

Likelihood ratio tests, on subsets I and II, are used to test for robustness. We estimate parameters for one subset of games by maximum likelihood. We then impose these parameter estimates on a different subset and obtain the likelihood of the latter subset, called the *predicted subset*, with the imposed parameters.

We find that the likelihood-ratio statistic of subset I relative to the full set of players is 3.278, and the likelihood ratio of subset II relative to the full set is 4.776 (the 5% critical value is 12.59, χ^2 distribution with six degrees of freedom); the respective p-values are 0.773 and 0.573. Thus, we cannot reject the null hypothesis that the parameter estimates from the full set of games are valid for subsets I and II. Further, the likelihood-ratio statistic for predicting subset II from subset I is 9.849, with a p-value of 0.131.[18] Thus, we cannot reject the null hypothesis that the parameter estimates from subset I are valid for subset II. In other words, the parameter estimates are stable across these subsets of games, demonstrating the out-of-sample predictive power of the model.

6.2. Discussion

We wished to assess the relative success of various models of non-pecuniary motives in explaining this experimental data. First, it should be noted that each of the three approaches (altruism, distributive concerns, and intentional reciprocity) demonstrated substantial success in explaining the observed deviations relative to the null hypothesis of pecuniary considerations alone. Which is best? Can we gain much from combining their insights?

The so-called *equal division social reference point* of FS and BO seems to outperform a pure altruism explanation when employers determine wage. However, this reference point is inferior to pure altruism when wages are determined exogenously. This would seem to indicate that whatever equity considerations exist, they either (1) have a social reference point different than equal division,[19] or (2) a disutility from inequitable outcomes that is non-linear in distance from the reference point, or both. Overall, it seems that despite the important insight they provide, neither the simple altruism nor the pure equity-based models can explain behavior in the data analyzed in this essay.

The Reciprocity I model, which does not explicitly address distributive concerns, is also unsuccessful. On the other hand, the Reciprocity II formulation, which combines notions of reciprocity and altruism, does fairly well. However, estimation on the non-volitional treatments suggests some weaknesses in this approach. In particular, the success of the reciprocity term in that model seems to be due to capturing distributive concerns in the data rather than intentional reciprocity.

We also find that the restrictive equity formulation is rather successful among restrictive models, thereby reinforcing the assertion made in the last paragraph as well as suggesting that kindness may be measured in terms of actions rather than in terms of payoffs.

We find that among all models the unrestricted encompassing model performs best. Given the non-restrictive nature of the flexible approach, we can compare regard between the volition and non-volition treatments. The difference in regard can only be attributed to the difference in treatments; i.e., intention. Figure 3 displays regard as a function of wage for each treatment. Figure 4 shows the

[18] For subset I relative to the full set of players we get $-2[(-100.745) - (-99.106)]$; for subset II relative to the full set this is $-2[(-80.681) - (-78.293)]$. The likelihood-ratio calculation for predicting subset II from subset I is $-2[-42.0312 - (-37.1074)]$.

[19] Note that in real labor markets, the firm is making far more than the employee in monetary terms.

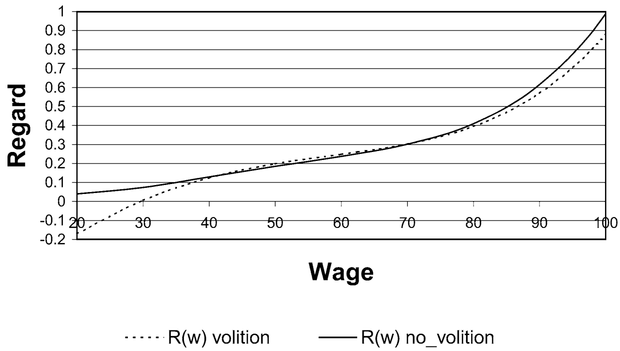

Fig. 3. Regard plotted on wage.

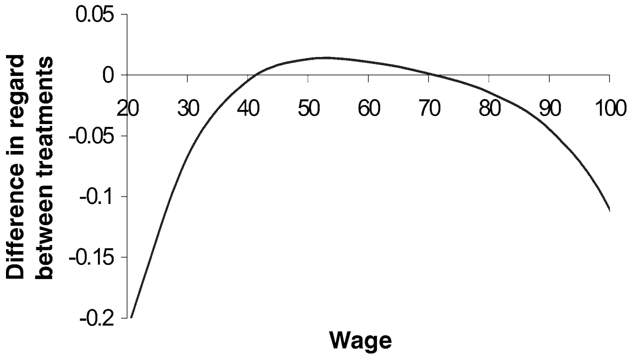

Fig. 4. Reciprocity (as measured by the difference in regard between treatments) plotted on wage.

difference in regard between treatments plotted on wage. Note that, in the center of the wage range, regard is roughly identical over treatments. The differences occur for “low” (roughly, wages less than 40) and “high” (roughly, wages greater than 85) wages. Not surprisingly, regard is lower when the “low” wage can be attributed to unkind, or selfish, intentions on the part of the employer. A surprising finding is that regard is lower in the volition treatment for “high” wages. One possible explanation is that when an outcome “unfair” to one party is willingly initiated by that party, the other party does not feel an obligation to reduce this inequality. On the other hand, if a provisional outcome is unfair to a party unable to participate to that point, the other party may feel an obligation to remedy the injustice.

7. Conclusion

As non-pecuniary motives are becoming recognized in mainstream economics and models of non-pecuniary behavior are proposed, it is useful to make comparisons of the effectiveness of such models in explaining observed behavior. We do so by nesting a number of models of pecuniary and non-pecuniary behavior within each other as well as within a flexible encompassing approach. Nested comparisons reveal that distributive concerns, altruism, and intentional reciprocity each contribute significantly to explaining non-pecuniary motives. Furthermore, comparing FS to a flexible approach and Reciprocity II in the non-volition treatments seems to suggest that distributive concerns are not likely to be linear in the payoff difference.

Another contribution of this essay is the flexible approach to estimation and testing. While this approach cannot be generalized to all settings, it is nonetheless warranted where feasible. Though parsimony is reduced by adding parameters, we believe it is valuable to combine the insights of the various approaches into a richer model. While we do not offer a complete model or a theoretical justification for the form of our flexible regard function or for the equity formulation, it is clear that these do a better job of fitting the data.

The flexible formulation (in particular its nesting properties and good fit of the data) allows us to separate intentional reciprocity from distributive concerns, demonstrating the existence of reciprocal behavior at low and high wages. We find that while unkind behavior is reciprocated with unkind behavior, kind behavior by the employer that puts the employer at a disadvantage may result in less attention to equity by the employee, as the unequal outcome is not considered to be the responsibility of the employee.

We feel that any successful model must accommodate the concerns of altruism, distributive concerns, and intentional reciprocity and suggest that further research is needed to develop a more descriptive model. It is true that our conclusions must be limited, as we only analyze data from one experiment. However, while we offer no formal analysis, evidence from some recent experiments also supports the view that intentional reciprocity plays an important role in non-pecuniary motives. Abbink et al. (2000) and Offerman (1998) show strong effects for reciprocity and the latter paper also indicates that the causal attribution for an outcome significantly affects the sequential choice made. These results point to the need for alternative explanations of observed non-pecuniary behavior.

It may be presumptuous to even suggest that one can identify all the important influences on choices made in economic contexts. Yet we hope that analyses such as the one in this essay can lead to descriptive models which both capture important motives and preserve substantial parsimony.

Acknowledgments

We thank Antonio Cabrales, Li Gan, Matthew Rabin, and Dale Stahl for helpful comments, and Kitt Carpenter and Chris Meissner for valuable research assistance. Financial support from Spanish Ministry Grant D101-7715 and the MacArthur Foundation is gratefully acknowledged.

Appendix A. Wage/effort pairs

Table A.1

	Effort									
Wage	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
	Employer-generated wages									
20	30									
21	5									
22	1									
25	5									
30	6									
35	1		1							
37		1								
40	7	13	4	1						1
45		3								
50	7	1	7	4						1
55	2			3						
60	9	3	2	10	6	5				
61				1						
65	1	1	1	1			1			
70	6		2	1	5	5	6	2	2	
75	3		1	2			1	1		
79	1						1			
80	5		1		2		2		1	2
85	3									
90	3						2			1
92	1									
94	1									
95			1				1	1		
97	1									
98										1
99								1		
100	2						1			1
101										1
105							1			

Table A.1 (*Continued*)

Wage	Effort									
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Random-generated wages										
20	17	3	2							
25	2		1	1						1
30	2									
35	2	2								
40	9	1	3							2
45	5	2			1					
50	4	1	2	3		1				2
55		1		1					1	
60	10	4	7	10	5		1		1	3
65	1		2							1
70	8	4	1	2	7		4	1		1
75	4	3		1	2	3		2		
80	4	1	4	3	1	1	4	1	1	
85						2	1	1		
90	2					1		2	1	
Third-party-generated wages										
20	14		1							
25	3		2							
30	1	2								
35	7		1		1					
40	9	3	4	1						
45	1	2	3	2						
50	3		4	5			2			
55	2			1						
60	13	2	4	7	14	2	1			
65			2	1		1				
70	10	2	2	5	10	2			2	
75	2	2	2	1	1	2	3			
80	10			1	2	2	7			1
85	1			1		1	1	1		
90	1			1		2	1			

Appendix B. Time effects

If reputation were important in effort determination, we should expect effort, or the effort-wage ratio, computed $(\text{effort} - 0.1)/(\text{wage} - 20)$, to decrease over time.

A casual look at the average wages and effort levels does not seem to support learning taking place in any period. The average-wage line seems relatively flat and the average effort level does not seem do display a pattern over time.

Visual evidence must be supplemented by rigorous testing. Let $\bar{q}_t$ be the average effort at time t. The F-test is a joint test of the hypotheses: Ho: $\bar{q}_{10-\delta+1} =$

$\cdots = \bar{q}_{10}$, which can take into account individual effects. Rejection of H_0 is necessary and sufficient to reject pooling.

We briefly describe the F-test: The F-test compares the unrestricted regression, $q_{it} = \mu + \alpha_i d_i + \beta_t D_t + \varepsilon$, to the restricted regression, $q_{it} = \mu + \alpha_i d_i + \varepsilon$, where d_i is an individual dummy, D_t is a time dummy, and α_i, β_t, and μ are parameters to be estimated. We impose two identification restrictions: $\sum_i \alpha_i = 0$

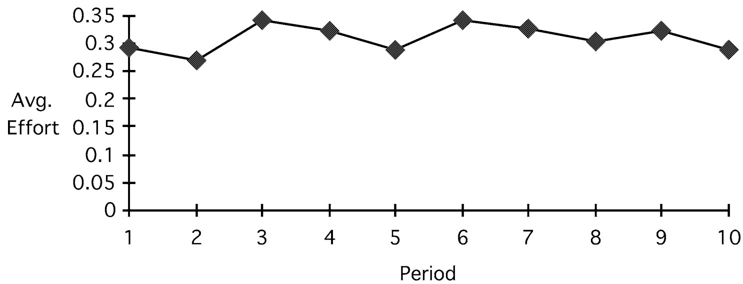

Fig. 5. Average effort by period

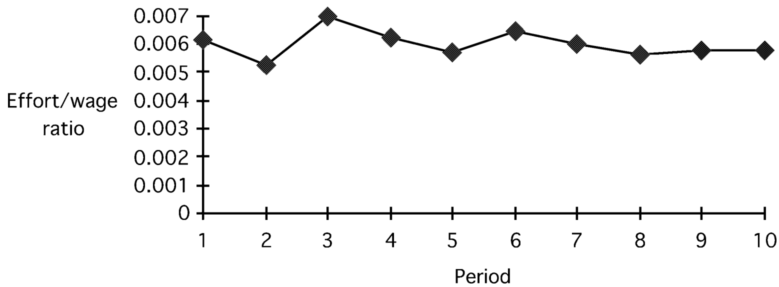

Fig. 6. Effort/wage ratio by period

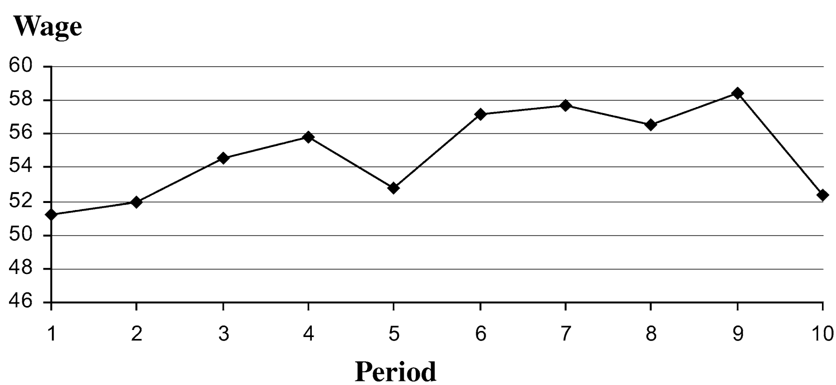

Fig. 7. Average wage by period

Table B.1
Time trends

Period	Avg. wage intentional	Avg. effort intentional	(Effort −0.1)/ (wage −20) intentional	Avg. wage non-intentional	Avg. effort non-intentional	(Effort −0.1)/ (wage −20) non-intentional
1	51.2	0.2909	0.00612	58.1	0.3513	0.00660
2	52.0	0.2682	0.00526	57.6	0.3333	0.00621
3	54.6	0.3409	0.00697	57.3	0.3180	0.00584
4	55.8	0.3227	0.00623	56.8	0.3513	0.00683
5	52.8	0.2864	0.00569	57.7	0.3231	0.00592
6	57.2	0.3409	0.00647	55.9	0.3000	0.00557
7	57.7	0.3273	0.00603	56.7	0.3154	0.00587
8	56.6	0.3046	0.00559	53.3	0.3359	0.00708
9	58.4	0.3227	0.00580	56.7	0.3026	0.00552
10	52.4	0.2864	0.00575	64.4	0.3744	0.00618

and $\sum_t \beta_t = 0$. It is easily verifiable, given the restrictions, that the estimate for μ will be the mean effort.

For each run, we first calculate the mean effort for the entire population over all time periods, $\bar{q}$, the mean effort for each individual i over all time periods, $\bar{q}_i$, and the mean effort in each time period t over all individuals, $\bar{q}_t$. The estimated individual effect (subject to the identification restrictions) for individual i is $\hat{\alpha}_i = \bar{q}_i - \bar{q}$. The estimated time effect in period t (subject to the identification restrictions) is $\hat{\beta}_t = \bar{q}_t - \bar{q}$. The F-test is as follows:

(1) We compute the unrestricted residuals by subtracting from each observation on effort, q_{it}, the average effort, $\bar{q}$, the estimated individual effect, $\hat{\alpha}_i$, and the estimated time effect, β_t.
(2) We compute the restricted residuals by subtracting from each observation on effort only the mean effort and estimated individual effect.
(3) The F-statistic is calculated as $[(\text{RSSR} - \text{USSR})/(T-1)]/[\text{USSR}/(NT - (N-1) - (T-1) - 1)]$, where RSSR is the restricted sum of squared residuals, USSR is the unrestricted sum of squared residuals, T is the number of time periods, and N is the number of individuals.

The F-statistics for the treatments of employer-determined wages, bingo-cage wages, and experimenter-determined wages are 0.254, 0.777, and 0.300, respectively. Under the null hypothesis of no time effects, these statistics should be distributed F (9, 188), F (9, 161), and F (9, 170), respectively. The respective critical values at the 5% level of significance are 1.930, 1.938, and 1.935. Hence we cannot reject the null hypothesis of no time effects at the 5% significance level.

References

Abbink, K., Irlenbusch, B., Renner, E., 2000. The moonlighting game—An experimental study on reciprocity and retribution. J. Econ. Behav. Org. 42, 265–277.

Andreoni, J., Brown, P., Vesterlund, L., 1999. What makes an allocation fair? Some experimental evidence. SSRI Working paper 9904, University of Wisconsin-Madison.

Blount, S., 1995. When social outcomes aren't fair: The effect of causal attributions on preferences. Org. Behav. Human Dec. Process. 63, 131–144.

Bolton, G., 1991. A comparative model of bargaining: Theory and evidence. Amer. Econ. Rev. 81, 1096–1136.

Bolton, G., Ockenfels, A., 2000. ERC: A theory of equity, reciprocity, and competition. Amer. Econ. Rev. 90, 166–193.

Brandts, J., Charness, G., 1999. Punishment and reward in a cheap-talk game. Mimeo.

Charness, G., 1996. Attribution and reciprocity in a simulated labor market: An experimental investigation. Mimeo.

Charness, G., Rabin, M., 2001. Some simple tests of social preferences. Quart. J. Econ., forthcoming.

Costa-Gomes, M., Zauner, K., 2001. Bargaining behavior in Israel, Japan, Slovenia, and the United States: A social utility analysis. Games Econ. Behav. 34, 238–269.

Dufwenberg, M., Kirchsteiger, G., 1998. A theory of sequential reciprocity. Working paper 1998:1, Stockholm University Department of Economics.

Fehr, E., Kirchler, E., Weichbold, A., Gachter, S., 1998. When social forces overpower competition: Gift-exchange in experimental labor markets. J. Labor Econ. 16, 324–351.

Fehr, E., Kirchsteiger, G., Riedl, A., 1993. Does fairness prevent market clearing? An experimental investigation. Quart. J. Econ. 108, 437–459.

Fehr, E., Schmidt, K., 1999. A theory of fairness, competition, and cooperation. Quart. J. Econ. 114, 817–868.

Goranson, R., Berkowitz, L., 1966. Reciprocity and responsibility reactions to prior help. J. Person. Soc. Psych. 3, 227–232.

Greenberg, M., Frisch, D., 1972. Effect of intentionality on willingness to return a favor. J. Exper. Soc. Psych. 8, 99–111.

Güth, W., Tietz, R., 1990. Ultimatum bargaining behavior: A survey and comparison of experimental results. J. Econ. Psych. 11, 417–419.

Heider, F., 1958. The Psychology of Interpersonal Relations. Wiley, New York.

Kagel, J., Kim, C., Moser, D., 1996. Fairness in ultimatum games with asymmetric information and asymmetric payoffs. Games Econ. Behav. 13, 100–110.

Kahn, A., Tice, T., 1973. Returning a favor and retaliating harm: The effects of stated intentions and actual behavior. J. Exper. Soc. Psych. 9, 43–56.

Kahneman, D., Knetsch, J., Thaler, R., 1986. Fairness and the assumptions of economics. J. Business 59, S285–300.

Kelley, H., Stahelski, A., 1970. Errors in perception of intentions in a mixed-motive game. J. Exper. Soc. Psych. 6, 379–400.

Ledyard, J., 1995. Public Goods: A Survey of Experimental Research, in: Kagel, J., Roth, A. (Eds.), Handbook of Experimental Economics. Princeton University Press, Princeton, pp. 111–194.

Levine, D.K., 1998. Modeling altruism and spitefulness in experiments. Rev. Econ. Dyn. 1, 593–622.

Ochs, J., Roth, A., 1989. An experimental study of sequential bargaining. Amer. Econ. Rev. 79, 354–384.

Offerman, T., 1998. Hurting hurts more than helping helps: The role of the self-serving bias. Mimeo.

Rabin, M., 1993. Incorporating fairness into game theory and economics. Amer. Econ. Rev. 83, 1281–1302.

Roth, A., 1995. Bargaining experiments, in: Kagel, J., Roth, A. (Eds.), Handbook of Experimental Economics. Princeton University Press, Princeton, New Jersey, pp. 253–348.

Thibaut, J., Riecken, H., 1955. Some determinants and consequences of the perception of social causality. J. Personality 24, 113–133.

Thomas, G., Batson, D., 1981. Effect of helping under normative pressure on self-perceived altruism. Soc. Psych. Quart. 44, 127–131.