

IZA DP No. 4084

**Do the Selfish Mimic Cooperators?
Experimental Evidence from Finitely-Repeated
Labor Markets**

Brian E. Roe
Steven Y. Wu

March 2009

Do the Selfish Mimic Cooperators? Experimental Evidence from Finitely-Repeated Labor Markets

Brian E. Roe

Ohio State University

Steven Y. Wu

*Purdue University
and IZA*

Discussion Paper No. 4084

March 2009

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Do the Selfish Mimic Cooperators? Experimental Evidence from Finitely-Repeated Labor Markets^{*}

Experimental studies have consistently shown that cooperative outcomes can emerge even in finitely repeated games. Such outcomes are justified by existing reputation building models, which suggest that cooperative outcomes can be sustained if some subjects have other-regarding preferences. While the existence of other-regarding preferences is typically used to justify experimental outcomes, we are unaware of empirical studies that explicitly examine the interaction between *cooperators* (those with other-regarding preferences) and selfish subjects in sustaining cooperation. In this paper, we classify subjects as either selfish or cooperative using simple social preference games and then test for behavioral differences between the two types in a finitely-repeated labor market with unenforceable worker effort. Theory predicts, and our data confirms, that (1) selfish players mimic the actions of cooperators when trading partners can track the individual reputation of past partners and (2) selfish and cooperative types act differently when individual reputations cannot be tracked.

JEL Classification: C91, D31, D86, K12

Keywords: contracts, relational contracts, implicit contracts, market interaction, experimental economics, repeated transaction, social preferences, reputation, firm latitude, finitely-repeated games

Corresponding author:

Steven Y. Wu
Department of Agricultural Economics
Purdue University
403 West State Street
Krannert Bldg.
West Lafayette, IN 47906
USA
E-mail: sywu@purdue.edu

^{*} The authors acknowledge helpful feedback from Ernst Fehr, John Kagel, Lorenz Goette and seminar audiences at the University of Toulouse, INRA-Paris, University of Verona, Ohio State University, Washington State University and the University of Zurich and research assistance from Myoungki Lee, Yun Jae Hwang, and Michael Brady. The authors also thank Martin Brown, Armin Falk, and Ernst Fehr for supplying detailed information concerning their experiments and associated Z-Tree programming which facilitated implementation of our experiments. The authors acknowledge support from the Ohio Agricultural Research and Development Center and the Cooperative State Research, Education and Extension Service, U.S. Department of Agriculture, National Research Initiative, Markets and Trade Division under Agreement Nos. 2003-35400-12887 and 2005-35400-15963. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the view of the U.S. Department of Agriculture.

Do the Selfish Mimic Cooperators?

Experimental Evidence from Finitely-Repeated Labor Markets

1. Introduction

Both theory and experimental evidence has shown that cooperative outcomes can arise in finitely repeated games where individual actions can be linked to agents' identity (e.g. David M. Kreps *et al.* 1982) and in finitely repeated games where interactions are anonymous (Healy 2007). The typical justification for cooperation in these environments is that if it is common knowledge that a certain fraction of agents are "cooperative" in that they are fair-minded and exhibit other-regarding preferences, then cooperative outcomes can emerge. Intuitively, in environments described by Kreps *et al.* cooperative agents will honor agreements even in the final round of a finitely repeated game, which then creates the potential for surplus that can be used to entice cooperation in earlier rounds. This then inspires even selfish agents to mimic the behavior of cooperative types because selfish agents can benefit from establishing a reputation for being cooperative and accessing final-period surplus that is available to cooperators.

Even in finitely repeated games with anonymous interaction, Healy (2007) has shown that cooperation is possible when group reputation is at stake. For example, if firms believe that workers' types are correlated, and workers know this, then even selfish workers may think twice before shirking as a single defect might damage workers' collective reputation to the point where only low wages are offered in the future. In a nutshell, the key mechanisms driving cooperation in finitely repeated games is (1) a sufficient belief in the existence of cooperative types and (2) the incentive for selfish types to mimic the behavior of cooperative types in all but the final period of a finitely repeated game.

While a number of theoretical and experimental studies have appealed to the mimicking argument to justify cooperative behavior in finitely repeated games, we are unaware of studies that explicitly examine the nature of interactions between selfish and cooperative types and the degree to which mimicking occurs in finitely repeated games. In this paper we use a series of abstract one-shot social-preference games to classify each experimental subject as either selfish or cooperative and then test for differences in the behavior of these two types of subjects in a finitely-repeated labor market experiment with unenforceable worker effort.

We exploit a between-subjects design in which different treatment groups are associated with distinct predictions regarding mimicking behavior during early rounds of play and during the final round of play. Some subjects participate in experimental labor markets that feature individual reputation tracking, which sharpens the incentives for a selfish second mover to mimic a cooperative second mover during early trading periods before reverting to type in the final round. The remaining subjects participate in labor markets that do not allow for individual reputation tracking, which dampens the incentive for a selfish subject to feign cooperation. We also consider two forms of labor markets: one in which workers move last in each period after receiving their wage and one in which firms move last in each period after observing worker effort.

Specifically, our experimental design includes four treatments. In our first treatment each firm offers a contract that promises a price P to workers and requests that workers exert some desired effort level. If a worker accepts a contract, she knows that she will be paid a guaranteed P , which does not depend on the observable but unenforceable effort she exerts. However, all subjects keep the same ID numbers through all periods of the experiments so that the formation of personal reputation can serve as a contract enforcement mechanism. We call this the “reputation” (R) treatment, which is similar to a design used by Brown, Falk and Fehr (2004).

Our second treatment is identical to the first treatment except ID numbers are randomly re-assigned after each period. Hence, firms and workers cannot identify each other across rounds so that there is only anonymous interaction. Only group reputation proposed by Healy (2007) can serve as a contract enforcement mechanism. We call this the “group reputation” (GR) treatment.

Our third and fourth treatments are similar to the first two, only now the last mover in each period is the firm rather than the worker. The third treatment is identical to the (R) treatment except that after the player in the role of the firm observes the effort level, he may make discretionary price adjustments; that is, he can refuse to pay, pay a reduced wage, or even pay more than the originally promised wage. We call this the “reputation with latitude” (RL) treatment. Finally, the fourth treatment is identical to (RL) except all ID numbers are randomly reassigned across periods. We call this the “group reputation with latitude” (GRL) treatment.

The experimental results regarding mimicking largely adhere to our predictions. In the reputation (R) treatment, after a brief learning period in which selfish workers act to type (exert less effort), selfish workers mimic the actions of cooperative types until the final round of play. In the final round of play, selfish workers once again exert less effort than do cooperative workers facing similar wages. In sessions where only group reputation is possible (GR) selfish workers exert less effort in every period.

When the order of play within a period is switched such that the firm moves last, we again find that the presence of individual reputation causes selfish second movers – firms in this case – to mimic cooperative types during all but the final period. When individual reputations cannot be tracked, however, the price a selfish firm offers in response to increased worker effort is less responsive than the prices offered by cooperative firms.

Our results suggest that, across two distinct market structures, selfish types quickly recognize that it is beneficial to imitate the behavior of cooperative types and then act upon this knowledge in the context of the labor market. An individual reputation mechanism leverages the benefits that a cooperative type confers to his or her individual trading partner across the entire subject population by inducing selfish types to act in a cooperative fashion during most trading periods. This cooperation leads to higher levels of effort which, for the markets we consider, translates to larger social surplus. When individual reputation can no longer be tracked, incentives for mimicking are significantly weakened. Indeed, the behavior of the players we classify as selfish becomes statistically distinct and economically burdensome on the entire market, suggesting that our classification system has predictive power in settings beyond the simple social preference games used to sort players into ‘selfish’ and ‘cooperative’ groups.

Our work is closest to that of Fischbacher and Gächter (2008), who also combine within-subject and between-subject design elements, though they establish that selfish and cooperative types act differently in a public goods game with random re-matching.¹ Page, Putterman and Unel (2005) find selfish subjects mimicking the actions of cooperative types in a repeated public goods game with endogenous re-matching based on past actions, though they focus on the context of public-goods games and they rely upon a subject’s final-period behavior within the same public goods experiment as the basis for assigning subjects to a social preference class.

¹ They also rely on a between-subjects comparison to verify the exogeneity and consistency of player type, whereas we rely upon a test-retest within-subject approach.

The remainder of the paper is organized as follows. We introduce the experimental design in section 2 and review the results of the social preference classification experiments in section 3. We then derive some predictions for the labor contracting experiment in section 4 and discuss the results of these in section 5. Section 6 concludes.

2. Experimental Design

All subjects participate in two experiments: one to assign them to a social preference (SP) type and one to test for behavior in a labor market setting with unenforceable worker effort. The flow of the experimental session and experimental design are outlined in figure 1.

Subjects' social preferences are classified by implementing a sequence of eight games adapted from Charness and Rabin (2002) and by implementing a classification scheme inspired by Andreoni and Vesterlund (2001). We pair subjects together and require each subject to make decisions that affect the monetary payoffs of both pair members. Each game allows subjects to alter the pair member's payment, usually by forgoing own payment (an example is included in the appendix containing experimental instructions and a summary of the games is provided in Table 1). By having each subject participate in a sequence of games with variations in roles and possible outcomes, various social preferences can be identified.

We categorize subjects as either selfish or cooperative. Under the cooperative class, we include individuals that exhibit disadvantaged inequality aversion – ensuring payments are as equal as possible so long as own payment is less than the payment received by their pair member; maximin preferences – choosing the option where the least-well-off player's payment is highest; and efficiency preferences – maximizing the sum of own and pair member's payments. The utility functions for each preference class and the classification scheme are detailed in the appendix.

This protocol provides an intentional or 'cold' measure of social preferences, i.e., when the subject is a second mover in a particular game, the decision impacting the first mover is made prior to the revelation of the first-mover's action (i.e., the strategy vector approach), without

knowledge of the first-mover's identity, and with no opportunity to identify the first mover on subsequent occasions during any part of the experiment.²

The second experiment is a contractual labor market where effort exerted by workers is observable but not third-party enforceable. Subjects are randomly partitioned into two groups: firms and workers. Firms offer contracts to workers specifying a price-effort combination for the period.³ An offer can be made to all workers simultaneously with public knowledge that all other workers are receiving the exact same offer (public offer) or to a single worker with knowledge that no other worker is receiving the offer (private offer). A firm can submit as many public and private offers as desired. The first worker to accept the firm's offer is matched with that firm and both firm and worker exit the market for that period.

Workers outnumber firms, which imposes involuntary unemployment. Workers not participating in the market in any given period for any reason receive an unemployment benefit (or reservation wage, r). Furthermore, the worker can always shirk, i.e., exert effort less than that prescribed in the offer.

The firm's market parameters are:

$$(1) \quad \pi_f = \begin{cases} 10E - P & \text{if trade occurs} \\ 0 & \text{if no trade occurs} \end{cases}$$

where π_f is the firm's payment, E is effort exerted by the worker, and P is the agreed upon payment or wage. Worker profits are:

$$(2) \quad \pi_w = \begin{cases} P - c(E) & \text{if trade occurs} \\ r & \text{if no trade occurs} \end{cases}$$

where $c(Q)$ is a strictly increasing and convex cost function fully represented by the following schedule of quality-cost combinations: $\{1,0\}$, $\{2,1\}$, $\{3,2\}$, $\{4,4\}$, $\{5,6\}$, $\{6,8\}$, $\{7,10\}$, $\{8,12\}$, $\{9,15\}$, $\{10, 18\}$. Agreed upon effort (E^*) and exerted effort (E) are allowed to be any integer

² 'Hot' social preference designs, i.e., those that allow for response to specific actions and remove anonymity may result in different measures of social preference parameters. Previous research presents a mixed picture of whether 'hot' and 'cold' settings provide different (Shafir and Tversky (1992), Croson (1999); Güth, Huck and Müller (2001); Burks, Carpenter and Verhoogen (2003)) or similar (Cason and Mui (1998), Brandts and Charness (2000), Oxoby and McLeish (2004)) measures of underlying social preferences.

³ Experimental directions use the terms buyer, seller, price and quality, though we will interchangeably discuss results using the terms firm, worker, wage and effort. Fehr, Klein and Schmidt (2007, pg. 141) find no statistically significant change in behavior in a one-shot labor market experiment similar to one of our treatments when the terms firm, worker and effort are substituted for buyer, seller and quality.

from one to ten while the price (P) can be any integer from zero to one hundred. Social surplus (W) is $10E - c(E) - r$, which is monotonically increasing in quality ($\partial W/\partial E = 10 - c'(E) > 0$) because marginal effort cost never exceeds 3. Social surplus is maximized only when the worker exerts maximum effort ($W=82$ for $E=10$) and minimized in the absence of trade ($-r$).

From this basic market framework, we implement a 2 x 2 between-subjects experimental design. The first design element is the firm's latitude concerning wages while the second element pertains to reputation. In the latitude condition, the firm is the last-mover in each round. The firm sets the wage paid to the worker (P) after perfectly observing the effort provided by the worker; this wage may deviate from the wage established in the initial offer (P^*) but must remain an integer from zero to one hundred. This makes the firm's initial wage and effort offer cheap talk in that neither element will be binding. One way to interpret this treatment would be the case where payment takes place after a job is complete or after delivery. In other words, the worker may work for the agent, issue an invoice and then the firm will decide whether to pay or not. If payment is made, it might be discounted if the firm claims low effort or the firm may "tip" the worker by paying a wage that exceeds promised price. The worker's reservation wage, r , is 10.

In the no latitude condition the worker is the last mover in each period. It is identical to the latitude condition, except that the firm is forced to pay the wage established in the initial offer regardless of the worker's effort ($P = P^*$). In other words, the payment is made before delivery or before work is completed. The worker's reservation wage, r , is 5. Note that both conditions feature unenforceable worker effort. The key difference is that, in the no latitude condition, the worker is the last to exert contractual latitude in a given round while in the latitude condition, the firm is the final mover.

The other design element is a manipulation of reputation. In the individual reputation condition, each party retains a unique identification number across all rounds and, at the end of each trading round, each subject is informed of the effort exerted, final wage and the payoffs for firm and worker. This form of reputation development does not facilitate a subject's global knowledge of all other subjects' past behavior, but does provide the ability to evaluate and act upon the past performance of previous partners. This structure of reputation tracking is at the heart of relational contracting, i.e., relationship-specific rents can be earned if parties grow to

trust one another over time while performance-contingent renewal, facilitated via private offers, can be wielded as a disciplinary ‘stick’ to foster that trust. In the group reputation condition, the identification numbers for both firm and worker are re-assigned after each round of trading, which makes the tracking of individual reputation impossible. However, subjects may display concern for the reputation of all members of its group (workers or firms).

We implement four treatments. The reputation (R) treatment involves no firm latitude and fixed ID numbers across rounds so that individual reputations can be established. This treatment is similar to the ICF treatment used by Brown, Falk and Fehr (2004) and identical to the IC1 treatment used by Wu and Roe (2007); it is also similar to the Repeat Purchase Mechanism first introduced by Klein and Leffler (1981). The group reputation (GR) treatment is identical to the (R) treatment except that ID numbers are randomly re-assigned each round some firms and workers remain anonymous to each other. Hence, individual reputation cannot be established, although group reputation may still matter. This treatment is similar to the ICR treatment used by Brown, Falk and Fehr (2004). The third treatment is the reputation with latitude (RL) treatment which fixes ID numbers across rounds and provides firms with discretionary latitude to make price adjustments after effort is observed. In this treatment, firms have two incentive devices for informally enforcing effort – discretionary price adjustments and individual reputation. This treatment is identical to the IC2 treatment used by Wu and Roe (2007). Finally, the group reputation with latitude (GRL) is identical to RL except ID numbers are re-assigned after each period so that firms and workers remain anonymous to each other. Hence, only group reputation and discretionary price adjustments can informally enforce effort. This treatment is unique to this paper.

Each contracting session involves two practice rounds and fifteen rounds that determine eventual cash payment where each round lasts 90 seconds and experimental profit was converted to subject payment at a rate of 70 experimental monetary units to \$1. Practice rounds feature only the bidding stage of contracting to avoid costless, deceptive signaling by workers and firms in treatments that allow reputation formation. Each session features five firms and seven workers who were randomly assigned to these roles and retained these roles during all rounds. All trading takes place on networked computers enclosed in cubicles to eliminate between-subject visual contact. The 216 subjects were students recruited via e-mail and newspaper from

various academic departments at the authors' home institution and provided informed consent to participate in the protocol, which was approved by the local institutional review board.⁴

For the SP games, each subject played both roles within each game. Payment was made for one role from one of the eight SP games played where the exact role and game was chosen via a publicly observed random draw at the end of all experiments; one dollar was paid per 100 experimental monetary units earned.⁵

Others have suggested that SP classification may be sensitive to the players' experience in other games within the same experimental session or may be unstable over time.⁶ We test for the stability of the SP classification scheme over time and for sensitivity of classification to the labor market structure encountered by the subjects during the labor market by dividing the SP games into two sets of four games with one set administered at the beginning of the evening and one administered after the labor market session. The order of the two sets and the order of games within the two sets were counterbalanced across subjects.

Several other authors utilize a within-subject design featuring simple social preference games. For example, Blanco, Engelmann and Normann (2006) require subjects to cast seven decisions during four simple games (ultimatum, dictator, public goods and sequential prisoner's dilemma games). The authors then use the continuous responses captured in the ultimatum game to calibrate parameters for a Fehr and Schmidt (1999) model of inequality aversion, which are then used to predict individual and aggregate behavior in the other three games. Brosig, Riechmann and Weimann (2007) implement several variants of dictator and sequential prisoner dilemma games during a given session and then repeat the session two times over the next three months in an attempt to test for consistency across games and stability of social preferences across time. Fishbacher and Gächter (2008) use simple one-shot games to assign subjects a

⁴ Business or economics majors comprised 49% had of the subjects, while 13% had a quantitative focus (e.g., engineering or math) and the rest had another major. Women comprised 37% of the subject pool while 73% self-identified as Caucasian.

⁵ This was done to minimize income effects among the multiple decisions cast. Subjects are fully informed of this compensation tactic at the beginning of the experiment and advised to treat each choice with equal seriousness.

⁶ For example, Brandts, Riedl and van Winden (2006) find evidence of such feedback in an experimental setting where social preferences are measured before and after exposure to different degrees of market competition. Also, Healy (2007) suggests allowing for time-varying classification of subjects after reflecting upon the inter-period heterogeneity observed within subjects in his data.

social preference type and then use this type to predict subject behavior in a subsequent repeated public-goods session.

Our design differs from these previous efforts in that we use within-subject design elements to assign subjects to social preference types and a between-subject design to verify if different types respond to different institutional elements in a manner consistent with full reputation equilibrium theory in finitely repeated labor markets. In contrast, Fishbacher and Gächter (2008) focus on public-goods scenarios where the role of social preferences is well established. Finally, unlike Brandts, Riedl and van Winden (2006) we are interested in using the SP classification obtained in the simple games to explain behavior in the more complex setting of a repeated labor market, whereas their primary interest is the impact of the labor market on social preferences and social disposition.

3. Social Preference Classification

In the course of participating in the eight social preference games described in Table 1, each subject cast thirteen decisions (three games were dictator games and featured a decision for only one role). We use the eight decisions cast by the last or only mover in each game as input into our classification scheme.⁷ Our subjects generated responses that were similar to those reported in Charness and Rabin (2002). Using data published by Charness and Rabin, we find five of the eight responses used in analysis yield frequencies that are not statistically different from Charness and Rabin's subjects at the 10-percent level or less using a two-sided Fisher's exact test. We next classify each subject's 'closest' social preference class by comparing observed behavior and behavior predicted by the utility functions associated with the four social preference classes (see the appendix for more details and Andreoni and Vesterlund 2001 for a similar classification tactic). We find 29 percent of subjects perfectly adhere to the predictions of one of the four social preference classes;⁸ the remaining subjects are assigned to the closest class.

⁷ As was the case with the Charness and Rabin (2002) protocol, we do not collect information concerning the intentions of first-movers, which limits the usefulness of these decisions in assigning subjects to social preference classes and leads us to use only the subject's last-mover decisions.

⁸ Andreoni and Vesterlund (2001) find that 44 percent of their sample exactly fit one of their three social preference categories, which is statistically different from the percent of exact matches in our data (two-sided Fisher's exact test, $p = 0.002$).

Seventy-two percent of our subjects are classified as selfish.⁹ Our proportion of self-interested types is higher than both Andreoni and Vesterlund (2001) and Andreoni, Castillo and Petrie (2003) who both found about half of their samples adhered to pure self interest. Cooperative types are evenly distributed between firms and workers (Fisher's exact $p = 0.28$), across experimental sessions (Fisher's exact $p = 0.28$) and across the interaction of firm/worker and experimental sessions (χ^2 , $p = 0.44$).

Because half of the SP games contributing to the classification are administered after exposure to the labor market treatment, we can classify each subject twice (once using the SP games conducted at the beginning of the session and once using the SP games conducted after the labor market) and test to see if a subject's classification is consistent over time and exogenous to the labor market treatment. A clear majority of subjects (68%) are consistent in type;¹⁰ 18% switch from selfish to cooperative and 14% switch from cooperative to selfish.

We are also interested to see if social preference classifications respond to the labor market treatment imposed upon the subject. Two analyses are consistent with stability of type to labor market treatment. First, a logit regression of preference type as a function of treatment, game administration order, firm/worker status, and subject demographics is estimated with robust standard errors clustered on experimental session. The likelihood ratio test for the joint significance of the labor market treatment dummies yields a test statistic of 4.30 ($\chi^2(3)$, $p = 0.23$). Second, a multinomial logit regression of the change in type (cooperative to selfish, selfish to cooperative, or consistent) as a function of treatment, game administration order, firm/worker status, and subject demographics is estimated with robust standard errors clustered on session. The likelihood ratio test for the joint significance of the labor market treatment dummies yields a test statistic of 6.86 ($\chi^2(6)$, $p = 0.33$).

4. Labor Market Predictions

Under the assumption that all actors are self-interested, and that this self-interest is common knowledge, sharp predictions emerge for behavior in all four treatments. Three key features of

⁹ Twenty-nine percent of subjects perfectly adhere to selfish behavior; for the remaining 43% it is their closest preference type. Fishbacher and Gächter (2008) find 23% of their subjects to perfectly adhere to free riding behavior.

¹⁰ Those who were perfectly adhered to a social preference type in the first four games were less likely to switch: (21% versus 44%, Fisher's exact p -value = 0.002).

these predictions are that the surplus generated (1) is zero for treatments featuring firm latitude because the market collapses, (2) equals the minimal positive surplus allowed for treatments featuring no firm latitude and (3) is independent of the presence of a reputation mechanism.

These predictions are driven by the assumption of self-interest by all parties and by the finite horizon of the interactions, which unleashes standard unraveling arguments via backward induction. For example, in markets with firm latitude, the worker knows a firm can pay a wage of zero as the last move of the round. Given firms are selfish, workers assume they will receive zero pay and, rather than enter the market and receive zero pay, will choose to avoid the market and collect the positive unemployment benefit.

Once common knowledge of universal self-interest is no longer assumed, higher levels of surplus are predicted. The core driver of these predictions of higher surplus levels was first illustrated by Kreps et al. (1982) and extended by Healy (2007): given that a certain fraction of the population is comprised of cooperative types, then even a selfish type, who has a personal or group reputation to protect, may find it personally beneficial to mimic cooperative types in all periods except the final period.

In our context, depending on the treatment, either individual or group reputation is a crucial component when generating predictions because, when coupled with the private-offer mechanism, reputation converts the labor market from a finite sequence of anonymous “one-shot” games into a finitely-repeated game that can be played between agents. The individual reputation mechanism provides the strongest incentives for selfish second-movers (workers in no latitude treatments and firms in latitude treatments) within a given round to mimic the actions of cooperative second-movers so that first-movers will believe they are coupled with a cooperative type and reciprocate with cooperative actions in the following round. When individual actions can be linked to individual identities, shirking can quickly lead to loss of a player’s reputation for being a cooperative type.

Prediction 1: If subjects believe that a non-trivial fraction of the population is comprised of cooperative types and subjects’ actions can be linked to their identities, selfish second movers will mimic the actions of cooperative second movers.

Such behavior is expected in all rounds except the final round. In the final round, selfish second movers no longer have an incentive to act cooperatively, and this knowledge will change the actions of the first mover. Selfish first movers who have renewed a trading relationship hope that their partner is truly cooperative and not just acting the part. In order to induce higher effort in case their partner is truly cooperative, the selfish first mover still acts generously. If the second mover is indeed cooperative, the selfish first mover gets higher material payouts. If the second mover is selfish instead, the selfish first mover gets lower payouts but suffers no additional disutility than the actual loss in payment.

A cooperative first mover in this same situation would suffer both the material loss in payment and additional disutility due to the resulting distribution of payments (i.e., the payments would be less equal, which lowers the utility of a cooperative type). Therefore, cooperative first movers are predicted to act less generously than selfish first movers in the final round.

Prediction 2: For continuing trading relationships, the final round actions by cooperative first movers will be less generous than selfish first movers.

Prediction 3: Selfish second movers no longer mimic the actions of cooperative second movers in the final round of treatments with individual reputation.

Even when individual reputation cannot be tracked so that only anonymous interaction occurs across periods, Healy (2007) suggest that selfish types may still mimic cooperative types in order to harvest group reputation. That is, if subjects believe that a non-trivial fraction of the population is cooperative *and* that types are correlated (i.e. stereotyping exists), then under certain payoff parameters, a selfish second-mover may still mimic a cooperative second mover in order to preserve the group's reputation. Nonetheless, the conditions that support cooperation under group reputation are more stringent than the conditions that support cooperation under individual reputation so that a selfish subject whose individual identity cannot be tracked across periods may find it more difficult to harvest the future stream of goodwill that might be generated from acting cooperatively. Therefore, selfish second movers may act to type and be

less generous. This diminishes surplus either directly, in the case of no firm latitude, by stimulating selfish workers to shirk, or indirectly, in the case of firm latitude, by discouraging workers to enter the market. When surplus is created, selfish second movers can act opportunistically and claim the surplus with lower future consequences, which increases the inequity of the distribution of surplus.

Prediction 4: When actions cannot be linked to individual identities selfish second movers will mimic the actions of cooperative second movers less frequently.

5. Labor Market Results

We first explore the beliefs held by subjects concerning the likely behavior of their trading partners as these beliefs are crucial precursors to any mimicking behavior between selfish and cooperative types.

Result 1. Subjects believe many other participants will respond cooperatively.

As part of the experimental protocol, firms are asked to predict the effort that will be provided by workers while workers predict the price firms will pay in contracts with firm latitude. From figure 2, we can confirm that, even in the very first period, only a small fraction of players in each treatment expect their counterpart to respond with full opportunism. That is, the first mover believes the second movers will meet surplus-generating actions with at least some cooperation. When firms hold price latitude, fewer than 10% of workers' expect complete opportunism regardless of the reputation regime or period. When workers move last, only 20% of firms believe workers will fully shirk during the first period of the group reputation (GR) treatment.

This result is critical because such beliefs are a necessary condition for mimicking behavior. If all subjects believed other participants to be fully selfish, then equilibrium outcomes would be either minimal (no firm latitude) or zero (firm latitude) as standard backward induction and unraveling intuition would predict.

Result 2. Selfish workers mimic cooperative workers in the reputation (R) treatment, but deliver lower effort under the group reputation (GR) treatment.

We verify this with the censored regression analysis of effort (table 2) where effort is modeled as a function of worker type, normalized price (price less the average price paid over all subjects in rounds in that treatment) and an interaction of these two terms. Separate models are estimated for both R and GR treatments.¹¹ Parameter stability tests (see footnotes for table 2) reveal the model parameters are stable over all rounds for the GR treatment (reported in the right-most columns of table 2), but that there is a behavioral shift after period 3 in the R treatment (left-most columns of table 2).

During the first few rounds, when individual reputation is tracked, selfish workers are predicted to exert one-third less effort than cooperative workers when both are paid the average price (4.61 versus 3.05 units¹², p -value < 0.001). The difference between the predicted effort of a cooperative and a selfish worker for a range of normalized prices is displayed as the dotted line in figure 3. However, for periods three through fourteen, selfish workers' effort choices are only 15 percent lower than cooperative workers' effort when average price is paid, which is statistically indistinguishable (6.14 versus 5.23 units, p -value = 0.185). This smaller difference in effort between selfish and cooperative workers is reflected by the dashed line in figure 3.

In the group reputation treatment selfish workers deliver average effort that is 73% lower than that delivered by cooperative types when average prices are offered (2.92 vs. 0.78 units, p -value = 0.017). Compared to the individual reputation treatment, the group reputation treatment resulted in the largest divergence in effort provision between selfish and cooperative types (solid line, figure 3). This result partially validates our social preference classification scheme by revealing that, compared to cooperative types, selfish workers exert less effort in order to enhance own payment, and that this tendency is exacerbated when personal reputation is not in jeopardy.

¹¹ Round 15 is excluded for the individual reputation treatment as theory predicts a different relationship between the explanatory variables and effort in the final round when individual reputation is tracked.

¹² Unless otherwise noted, all predicted values reflect uncensored effort levels.

However, the workers classified as selfish are not purely selfish because they respond to higher prices with higher effort. This can be seen in figure 4, where the predicted effort-response functions are upward sloping in price for both cooperative and selfish workers, even in the group reputation treatment. In fact, in both reputation treatments, the normalized price coefficient is positive and significant,¹³ while the interaction between selfish type and normalized price is not significantly different from zero, suggesting that both selfish and cooperative types reciprocate higher prices with greater effort.¹⁴ Selfish workers, however, tend to anchor this reciprocal effort from a lower ‘base’ level of effort, which is reflected in the regression coefficient associated with the ‘selfish’ classification and is reflected by the vertical space between the selfish and cooperative workers’ predicted effort response functions (figure 4).

Result 3. Selfish firms mimic the behavior of cooperative firms in the reputation with latitude (RL) treatment, but the wage payments made by selfish firms are less responsive to worker effort in the group reputation with latitude (GRL) treatment.

This result arises from the censored regression analysis of wages paid (summarized in table 3) where wage is modeled as a function of firm type, normalized worker effort (effort less the average effort across all employed workers in all rounds for that treatment) and an interaction of these terms. Separate models are estimated for each reputation treatment and both models reveal parameter instability, which necessitates separate regressions for early and late rounds (see table 3 footnotes).

When individual reputation is tracked, selfish firms mimic cooperative firms in both base price and in price response to delivered effort, regardless of the period. This is revealed by the insignificant coefficients on the selfish dummy and by the insignificant interaction terms in the first two columns of table 3.

¹³ There is one exception – period 15 for treatment R – which we discuss later.

¹⁴ We forward several possible explanations for this finding. First, the SP games allow only discrete choices while the contracting games allow for continuous responses, which may reveal more precisely the social preferences of the player. Furthermore, for the sake of tractability, our dichotomous classification scheme forced players into a category even when their actions during the social preference games were not perfectly consistent with selfish or cooperative behavior. Second, the context of the social preference and contracting games is distinct as players receive rapid feedback in the contracting sessions but not the social preference games, which could lead to differential behavior.

When only group reputation is possible, two differences emerge between selfish and cooperative firms. First, during all periods, selfish firms are less price-responsive to worker effort as is revealed by the negative, significant interaction terms for periods 1-12 and periods 13-15. For example, during periods 1-12, a cooperative firm is predicted to respond to worker effort that is one unit above average by increasing the worker's wage by 5.97 units¹⁵ while a selfish firm is predicted to increase worker wage by 28% less (4.27 units, p -value < 0.001). This difference in effort responsiveness spikes in rounds 13-15: a unit increase in effort by a worker results in a wage increase of 21.06 units by the cooperative firm and only 5.87 units by the selfish firm (p -value < 0.001).

The second difference between selfish and cooperative firms under the group reputation treatment emerges in periods 13-15. During these final periods selfish firms provide significantly higher prices than cooperative firms to workers that supply average effort (49.10 units more, p -value < 0.001). One possible explanation for this large price 'base' may stem from the trend in trading volume visible in figure 5. The volume of trade (solid line) in the GRL treatment diminishes throughout the course of the experiment. More than 90% of possible trades are completed during the first six rounds, but this begins to decay to the point where, by period 10, only 80% of possible trades are completed. Selfish types, who compose about 80% of trading firms in these rounds, may begin to realize that their past stinginess is driving down market participation via a group reputation mechanism and that this past behavior may preclude them from trading in subsequent periods. Indeed, from figure 5 we can see that prices paid, which steadily decline during the first six rounds of trading, begin to climb thereafter, which stabilizes trading volume until the final rounds.

Result 4. In the final round selfish first-movers are more generous than cooperative first-movers in trades that have an established trading history.

Recall, this prediction is driven by the assertion that cooperative types dislike unequal distributions of surplus so much that they are predicted to offer contracts that result in small but equal levels of surplus between the two parties. Selfish first-movers, on the other hand, continue

¹⁵ For ease of exposition we discuss the uncensored predictions unless otherwise noted.

their past generosity in an attempt to continue the cycle of cooperation previously displayed by the second mover with whom they have an established relationship.

Such a pattern is verified in the estimated regression parameters in table 4. For example, in the first column, we have the firm's final-round up-front wage payment in the no latitude treatments as a function of firm type, relationship length and the interaction of these two variables, where relationship length is assigned as zero for all trades in the group reputation treatment. For relationships that have persisted at least six rounds, selfish firms pay significantly higher up-front wages in the final round (Selfish + Selfish*Length | length = 6) = 8.84, $\chi^2 = 4.94$, $p = 0.046$). For contracts with firm latitude (second column), selfish workers supply significantly more effort than cooperative workers to firms with whom they have a relationship of any length (Selfish + Selfish*Length | length = 1) = 23.48, $\chi^2 = 26.66$, $p < 0.001$).

Result 5. In the final round a selfish second-mover's tendency to revert to type depends upon the treatment.

Recall, the prediction was that selfish types should revert to type, i.e., act selfishly, even if individual reputation enticed them to mimic cooperative behavior in earlier rounds. The regression evidence is mixed concerning this prediction. For contracts without firm latitude, selfish workers are less responsive to the firm's price in the final round when individual reputations are tracked (column 3, table 2, interaction term). For example, a cooperative worker will respond to a one unit increase in normalized wage an additional 0.39 units of effort while a selfish worker will only increase effort by 0.18 units (p -value = 0.060). However, selfish workers' base effort in this final period is not statistically distinct from cooperative workers (column 3, table 2, first coefficient, p -value = 0.398).

Firm behavior in the final round does not conform to our predictions, however. When individual reputations are tracked, selfish firms provide, on average, higher prices for average effort, though these insights are based upon a small number of observations including just three cooperative firms that traded in the final period. In fact, all three cooperative firms induced their workers to provide the maximum effort, which precludes estimation of an interaction term in the regression model of final period behavior.

6. Discussion and Conclusions

While imitation may be the sincerest form of flattery, we find that it is also the dominant mechanism used by selfish second movers within our experimental labor market. Selfish second-movers, who reveal selfish tendencies during simple, anonymous, cold social-preference games held before and after the experimental labor market, tend to mimic the behavior of cooperative second-movers when the market is structured in a way that allows players to individually track the reputation of those with whom they have previously interacted. These results are consistent with the theoretical predictions of well-known full reputation equilibrium models (e.g. Kreps, et. al. 1982; Healy, 2007)

This mimicking behavior is supportive of surplus in two distinct ways. First, in order to mimic cooperative second movers, selfish second movers must be more generous in their actions. In markets where the worker moves last, this means that selfish workers provide greater effort in the presence of individual reputation during all but the last period in order to disguise themselves as cooperative workers. This allows these workers to sustain long-term trading arrangements with firms that will then pay higher wages. For markets where firms move last, this means that selfish firms provide higher wages in the presence of individual reputation in order to disguise themselves as cooperative firms. This will allow the firm to sustain long-term agreements with workers who are willing to provide higher effort levels than could be obtained in the absence of reputation. In both cases the mimicking tendency sustains trades with high levels of effort which, in this labor market, translates to high levels of social surplus.

The second way that mimicking behavior is supportive of surplus is through its effect on the volume of trade. When selfish second-movers fail to mimic cooperative behaviors, and instead act in an opportunistic fashion, it can cause large losses for the first mover. After sustaining several rounds of large losses, some first-movers may become discouraged and leave the market altogether to collect the reservation payment, which causes no social surplus to be generated. Indeed, when mimicking is the norm (when individual reputation is tracked), more than 97% of all trades are transacted, regardless of whether firms move first or last. However, once individual reputation is not tracked, the percent of trades transacted by round 15 declines to less than 85%.

Mimicking behavior not only enhances surplus in our two experimental labor markets – it also has distributional implications. Specifically, surplus is more evenly distributed when mimicking predominates (individual reputation) than when selfish types fail to mimic. For example, in figure 8 we display the average surplus generated for each trading partner in each of our four treatments. It is clear that the first-mover in each treatment receives the smaller share of the surplus – firms in the case of no latitude treatments and workers in the case of firm latitude treatments. However, the figure also reveals that a switch from group reputation to individual reputation, which we have shown to induce selfish second-movers to mimic cooperative second movers, moves the distribution of surplus closer to a 50/50 split. For example, when firms move first, they receive about 25% of total surplus under group reputation but nearly 41% of surplus if individual reputations are tracked. Likewise, when workers move first, they receive only 25% of surplus with group reputation and about 35% with individual reputations.

The fact that individual reputation helps the second mover is quite intuitive. Individual reputation in a market with endogenous re-matching effectively converts this round's second mover into next round's first mover because this round's second mover can reject offers from an opportunistic player in subsequent rounds.

Our results concerning the interaction of selfish and cooperative types expand upon findings from public goods experiments. For example, both Page, Putterman and Unel (2005) and Fishbacher and Gächter (2008) find that group composition (i.e., the percent of free riders versus cooperators) affects the degree to which cooperation can be sustained in a public goods game. Particularly, Page Putterman and Unel (2005) find that with a repeated public goods game, cooperation is sustained with endogenous sorting of groups, i.e., when individual group members can choose future group members based on the observed past behavior of individual players. Much of the cooperation in their public goods games comes from mimicry, i.e., those who defect in the last round but choose to invest in public goods in earlier rounds. Our labor market structure, when coupled with individual reputation formation, also allows for endogenous re-matching of partners over periods, which also induces mimicry by the selfish for the good of the many.

Our finding of mimicking in the presence of individual reputation tracking is also reminiscent of patterns identified by Jin and Kato (2006) among baseball card traders on eBay.

Specifically, they find that some sellers act cooperatively (sell cards whose quality meets advertised levels) for a time in order to establish a good reputation. After establishing this reputation these sellers then misrepresent the quality of a number of goods in rapid succession in order to earn a high price for low-quality goods. These opportunistic transactions quickly depreciate their reputation and such sellers exit the eBay marketplace under that identity. These individuals may then re-enter the market at a later date under a new identity and repeat the sales cycle. As in our experimental labor market, the presence of reputation induces selfish types to mimic cooperative types until near the end of a finite horizon. The challenge for future research and institutional design is to construct and improve mechanisms such as reputation that can maximize the desire for selfish types to mimic cooperative behaviors for the benefit of all players in the market.

References

- Andreoni, James, March Castillo, and Raga Petrie. 2003. "What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game." *American Economic Review*, 93(3): 672-685.
- Andreoni, James and Lise Vesterlund. 2001. "Which is the Fair Sex? Gender Differences in Altruism." *Quarterly Journal of Economics*, 116(1):293-312.
- Blanco, Mariana, Dirk Engelmann, and Hans-Theo Normann. 2006. "A Within-Subject Analysis of Other-Regarding Preferences." Working paper, Royal Holloway, University of London.
- Brandts, Jordi and Gary Charness. 2000. "Hot versus Cold: Sequential Responses and Preference Stability in Experimental Games." *Experimental Economics*, 2:227-38.
- Brandts, Jordi, Arno Riedl and Frans van Winden. 2006. "Competition and Well-Being." Tinbergen Institute Discussion Paper, TI 2004-041/1, Tinbergen Institute, Rotterdam, Amsterdam.
- Brosig, Jeannette, Thomas Riechmann, and Joachim Weimann. 2007. "Selfish in the End?: An Investigation of Consistency and Stability of Individual Behavior." MPRA Paper #2035, University of Magdeberg, Magdeberg, Germany.
- Brown, Martin, Armin Falk, Armin, and Ernst Fehr. 2004. "Relational Contracts and the Nature of Market Interactions." *Econometrica*, 72(3):747-80.
- Burks, Stephen V., Jeffrey P. Carpenter, and Eric Verhoogen. 2003. "Playing Both Roles in the Trust Game." *Journal of Economic Organization and Behavior*, 51(2):195-216.
- Cason, Timothy N. and Vai-Lan Mui. 1998. "Social Influence in the Sequential Dictator Game." *Journal of Mathematical Psychology*, 42(2/3):248-65.
- Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics*, 117(3):817-69.
- Croson, Rachel. 1999. "The Disjunction Effect and Reason-based Choice in Games." *Organizational Behavior and Human Decision Processes*, 80(2):118-33.
- Fehr, Ernst and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics*, 114:817-868.

- Fehr, Ernst, Alexander Klein, and Klaus M. Schmidt. 2007. "Fairness and Contract Design." *Econometrica*, 75(1):121-54.
- Fischbacher, Urs and Simon Gächter. 2008. "Heterogeneous Social Preferences and the Dynamics of Free Riding in Public Goods." Thurgau Institute of Economics, Research Paper Series #27, May, Konstanz, Germany.
- Fudenberg, Drew and Eric Maskin. 1986. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica*, 54(3):533-54.
- Güth, Werner, Steffen Huck, and Wieland Müller. 2001. "The Relevance of Equal Splits in Ultimatum Games." *Games and Economic Behavior*, 37(1):161-69.
- Healy, Paul J. 2007. "Group Reputations, Stereotypes and Cooperation in a Repeated Labor Market," *American Economic Review*, 92(5):1751-1773.
- Jin Ginger and Andrew Kato. 2006. "Price, Quality and Reputation: Evidence from an Online Field Experiment," *RAND Journal of Economics* 37(4):983-1005.
- Klein, Benjamin and Keith B. Lefler. 1981. "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, 89:615-641.
- Kreps, David M., Paul Milgrom, John Roberts and Robert Wilson. 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory*, 27(2):245-252.
- MacLeod, W. Bentley and James M. Malcolmson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica*, 57:447-480.
- Oxoby, Robert J. and Kendra N. McLeish. 2004. "Sequential Decision and Strategy Vector Methods in Ultimatum Bargaining: Evidence of the Strength of Other-Regarding Behavior." *Economics Letters*, 84(3):399-405.
- Shafir, Eldar and Amos Tversky. 1992. "Thinking through Uncertainty: Nonconsequentialist Reasoning and Choice." *Cognitive Psychology*, 24(4):449-74.
- Wu, Steven Y. and Brian E. Roe. 2007. "Discretionary Latitude and Relational Contracting." IZA Discussion Paper #2879, Bonn, Germany, June.

Table 1. Subject Responses to Social Preference Games

Game	Role	Options	% Chose Left
1	A	No Choice	--
	B*	(400, 400) v. (750, 375)	80
2	A	(100, 1000) v. let B choose	--
	B*	(75, 125) v. (150, 125)	32
3	A	(700, 200) v. let B choose	--
	B*	(200, 700) v. (600, 600)	72
4	A	(375, 1000) v. let B choose	--
	B*	(400, 400) v. (350, 350)	95
5	A	No Choice	--
	B*	(300, 600) v. (700, 500)	79
6	A	(750, 0) v. let B choose	--
	B*	(400, 400) v. (750, 375)	78
7	A	(500, 500) v. let B choose	--
	B*	(800, 200) v. (0, 100)	79
8	A*	(550,550,550) v. (600,300, 900)	43
	B	No Choice	--
	C	No Choice	--

Notes: * Denotes decisions used to classify subjects into social preference categories. N=216

Table 2. Worker Effort in Contracts with No Firm Price Latitude.

	Individual Reputation			Group Reputation
	<i>Periods 1-3</i>	<i>Periods 4-14</i>	<i>Final Period</i>	<i>All Periods</i>
Selfish	-1.563*** (0.235)	-0.910 (0.686)	-1.945 (2.299)	-2.141** (0.893)
Normalized Offered Price	0.100*** (0.035)	0.204*** (0.046)	0.391*** (0.104)	0.15*** (0.024)
Selfish * N.O. Price	0.027 (0.058)	0.020 (0.053)	-0.210* (0.112)	0.031 (0.055)
Constant	4.612*** (0.538)	6.141*** (0.802)	4.300* (2.397)	2.925*** (0.842)
N	103	374	35	424
χ^2	53.34	37.08	19.53	85.20

Notes. Censored regression results with robust standard errors clustered on the experimental session (in parentheses). *,**,*** indicates significance at the 10, 5 and 1 percent levels. Individual reputation sample is divided into two subsamples (periods 1-3 and periods 4-14) based upon likelihood ratio testing ($\chi^2(4) = 38.94$, p -value < 0.001); alternative break points yielded smaller test statistics. Likelihood ratio testing did not justify splitting the Group Reputation sample at any point in time.

Table 3. Firm Price in Contracts with Firm Price Latitude.

	Individual Reputation			Group Reputation	
	<i>Periods 1-2</i>	<i>Periods 3-14</i>	<i>Final Period</i>	<i>Periods 1-12</i>	<i>Periods 13-15</i>
Selfish	-2.288 (6.316)	-1.782 (2.000)	19.141*** (4.840)	-0.442 (1.397)	49.103*** (4.460)
Normalized Effort	8.578*** (1.875)	5.871*** (0.381)	6.549 (4.628)	5.967*** (0.671)	21.059*** (0.823)
Selfish * N. Effort	-2.539 (2.035)	0.408 (0.456)	--	-1.703*** (0.521)	-15.193*** (0.548)
Constant	37.256*** (5.643)	32.240*** (1.032)	14.365 (13.651)	23.554*** (1.961)	-27.728*** (4.037)
N	40	233	18	320	66
χ^2	148.61	996.27	41.27	217.46	834.64

Censored regression results with robust standard errors clustered on the experimental session (in parentheses). *, **, *** indicates significance at the 10, 5 and 1 percent levels. The individual reputation sample is divided into three subsamples: periods 1-2, periods 3-14 and the final period. The split between periods 2 and 3 is based upon likelihood ratio testing ($\chi^2(4) = 9.75$, p -value = 0.045); alternative break points yielded smaller test statistics. Similarly, the group reputation sample is divided into two subsamples: periods 1-12 and periods 13-15. The split between periods 12 and 13 is based upon likelihood ratio testing ($\chi^2(4) = 14.52$, p -value = 0.006); alternative break points yielded smaller test statistics. For the individual reputation treatment, period 15 is analyzed separately because the theoretical predictions are qualitatively different than predictions for earlier periods. In the final-period all cooperative firms received maximum effort from their workers, which precludes the independent estimation of an interaction term between selfish type and normalized effort.

Table 4. First-Mover Behavior in the Last Period

	Firm's Wage No Latitude	Worker's Effort Latitude
Selfish	6.785 (7.336)	4.800 (3.791)
Relationship Length	2.978*** (0.549)	0.643 (0.473)
Selfish * Length	0.342 (0.847)	18.683*** (3.250)
Constant	15.134* (7.419)	6.644 (4.354)
Test: Selfish + Selfish*Length length = 1	1.16	26.66***
Test: Selfish + Selfish*Length length = 6	4.94**	37.02***
N	64	38
Goodness of Fit (R^2 and χ^2 , respectively)	0.22	43.22

The firm's wage model is a linear regression (no price observations were censored) with robust standard errors clustered on the experimental session (in parentheses); the test statistic is an F-statistic and the goodness of fit is R^2 . The worker's effort model is a censored regression with robust standard errors clustered on the experimental session (in parentheses); the test and the goodness of fit statistics are χ^2 . *, **, *** indicates significance at the 10, 5 and 1 percent levels.

<u>Sequence</u>	<u>Event</u>
Event 1	Social Preference Classification Randomly assign: Games 1-4 <i>or</i> Games 5-8
Event 2	Contracting Session Randomly assign: L-R: Latitude – Reputation <i>or</i> L-NR: Latitude – No Reputation <i>or</i> N-LR: No Latitude – Reputation <i>or</i> NL-NR: No Latitude – No Reputation
Event 3	Social Preference Classification Assign games not yet played: Games 5-8 <i>or</i> Games 1-4
Event 4	Collect Demographics and Payment

Figure 1. Experimental session design and sequence




Figure 2. Expected opportunism in first and later periods by treatment.




Figure 3. Predicted Effort Difference between Cooperative and Selfish Workers




Figure 4. Predicted Worker Effort (Uncensored) in Group Reputation Treatment




Figure 5. Trade Volume and Price Paid by Round in the Group Reputation with Latitude Treatment.




Figure 7. Surplus by Reputation and Firm Latitude by Period




Figure 8. Firm and Worker Surplus by Reputation and Latitude.

Appendix: Classifying Subjects into Social Preference Categories

Let $U_{j,g}^i(\pi_{j,g}, \pi'_{j,g})$ denote the utility derived by the subject for selecting the j th option in game g for the i th utility functional form, where $\pi_{j,g}$ denotes a subject's own payment for the j th option in game g and $\pi'_{j,g}$ denotes the payment of the subject's partner. We postulate four possible utility forms:

(Selfish)	$U_{j,g}^S = \pi_{j,g},$
(Disadvantaged Inequality Aversion)	$U_{j,g}^D = -\max[\pi'_{j,g} - \pi_{j,g}, 0],$
(Social Efficiency)	$U_{j,g}^E = \pi_{j,g} + \pi'_{j,g},$
(Maximin)	$U_{j,g}^M = \min[\pi_{j,g}, \pi'_{j,g}]$ and

During the social preference (SP) games, each subject faces the same eight decisions in the role of a last or only mover. In each decision the subject is provided two options, $j \in \{L, R\}$, where L denotes the left option and R the right option. For each utility function there exists a maximum (\bar{V}^i) and minimum (\underline{V}^i) utility obtainable across the eight games. For each subject and each functional form, utility is calculated for the choices made by the subject during the eight SP games, i.e., $V^i = \sum_{g=1}^8 U_{j,g}^i$. Normalized utility scores of the form $v^i = [V^i - \underline{V}^i] / [\bar{V}^i - \underline{V}^i]$ are then calculated for each subject and all possible utility forms. The subject is then assigned to the preference category that generates the highest normalized utility score. If a subject is assigned to a category other than selfish, the subject is classified as cooperative for the purposes of subsequent analyses.