Social Preferences and Relational Contracting Performance: An Experimental Investigation

Brian E. Roe and Steven Y. Wu*

Abstract. We examine how social preferences affect behavior and surplus in relational contracts. In our within-subject experimental design, subjects participate in a contracting environment similar to Brown, Falk, and Fehr [Brown, M., Falk, A. & Fehr, E., "Relational Contracts and the Nature of Market Interactions, *Econometrica,* 72 (2004):747-780] and in social preference experiments adapted from Charness and Rabin [Charness, G. & Rabin, M. "Understanding Social Preferences with Simple Tests." *The Quarterly Journal of Economics* 117(2002): 817-869]. Subjects' behavior during the Charness and Rabin experiment is a significant predictor of behavior and outcomes observed during the subsequent multi-period relational-contracting environment, which features buyer concentration, unenforceable performance, reputation formation and endogenous matching of trading partners. The nature of social preferences exhibited during the Charness and Rabin games manifests in a straightforward manner in the contractual setting, e.g., social efficiency sellers improve total trading surplus by shirking less while buyers, who typically earn less than sellers in these contracts, fare relatively well when either party exhibits maximin social preferences. We conclude that intentional or 'cold' measures of social preferences have considerable predictive power in dynamic, interactive (or 'hot') economic settings.

Key Words: Contracts; relational contracts; implicit contracts; market interaction; experimental economics; repeated transaction; social preferences.

JEL Codes: C91, D31, D86, K12

* Roe and Wu are Associate and Assistant professors, respectively, in the Department of Agricultural, Environmental and Development Economics, Ohio State University; senior authorship is shared equally. Roe and Wu acknowledge helpful feedback from John Kagel and Lorenz Goette and research assistance from Myoungki Lee, Yun Jae Hwang, and Michael Brady. Roe and Wu thank Martin Brown, Armin Falk, and Ernst Fehr for supplying us with detailed information concerning their experiments and associated Z-Tree programming which facilitated implementation of our experiments. Roe and Wu 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.

Social Preferences and Relational Contracting Performance: An Experimental Investigation

1. Introduction

In an effort to incorporate greater behavioral realism into economic models, researchers have developed and tested a growing array of social preference models, i.e., individual models of utility that deviate from models of pure self interest (Ernst Fehr and Urs Fischbacher (2002) provide an extensive review of modeling efforts; a partial list of more recent work includes Gary Charness and Matthew Rabin (2002), James C. Cox (2004), Martin Dufwenberg and Georg Kirchsteiger (2004), Dirk Engelmann and Martin Strobel (2004), Gary E. Bolton, Jordi Brandts, and Axel Ockenfels (2005), Bruno S. Frey and Alois Stutzer (2005), Yoella Bereby-Meyer and Muriel Niederle (2005), and Armin Falk and Fischbacher (2006)). Recent contributions in this field have focused on devising simple experimental games that allow for efficient testing of competing models of social or other-regarding preferences (e.g., Charness and Rabin (2002), Engelmann and Strobel (2004), Bolton, Brandts and Ockenfels (2005)).

The ultimate goal of such model development and testing is to predict how outcomes in more complex economic interactions will be affected by the presence of agents with social preferences. Significant experimental evidence documents outcomes that deviate from those predicted by models of pure self-interest (Fehr, Kirchsteiger and Arno Riedl (1993, 1996, 1998), Fehr, Simon Gachter, and Kirchsteiger (1997), Fehr et al. (1998), Fehr and Falk (1999), Charness and Ernan Haruvy (2002), R. Lynn Hannan, John H. Kagel, and Donald V. Moser (2002), Brandts and Charness (2004), Brandts, Riedl and Frans van Winden (2004), Charness (2004), Fehr and Klaus M. Schmidt (2004), and Martin Brown, Falk, and Fehr (2004)), and gauges the strength and nature of social preferences indirectly from the size and direction of the deviations from equilibrium values predicted under pure self interest. Less theoretical progress has been

made in developing solution concepts for richer, dynamic economic settings involving agents possessing social preferences, though several contributions are notable (Fehr and Schmidt (2000), Brown, Falk and Fehr (2004), and Dufwenberg and Kirchsteiger (2004)).

Previous theoretical work by Brown, Falk and Fehr (2004) establishes that the presence of enough agents with other-regarding preferences in a finitely repeated market game can generate equilibrium surplus higher than that obtainable when the market is populated with purely self-interested agents. Brown, Falk and Fehr (2004) also demonstrate that subjects participating in an experimental contracting market did indeed exceed the surplus predictions associated with a model featuring participants with purely self-regarding preferences.

Our contribution is to use a within-subject design to explicitly link a subject's social preferences as measured by simple, intentional social preference games to subsequent behavior in a complex, dynamic relational contracting setting. To facilitate comparability with previous contributions we utilize existing experimental protocol for measuring intentional social preferences (Charness and Rabin (2002), hereafter, CR) and for documenting behavior in relational contracting settings (Brown, Falk and Fehr (2004), hereafter BFF). The responses from the 'cold' CR setting are then used to explain behavior in subsequent 'hot' experimental market activity.

The BFF relational contracting environment provides a rich market landscape upon which to observe the influence of social preferences. Buyers initiate contracting offers and can extend offers to all sellers simultaneously (public offers) or to a single seller (private offers), which enables long-term relationships. Sellers outnumber buyers and are, therefore, subject to involuntary unemployment. However, sellers may unilaterally deviate from contract terms and claim substantial short-term rewards for such shirking. Buyers, while offered no immediate

recourse to such shirking, can track past partners' performance and avoid such sellers in subsequent dealings.

Our analysis of more than 500 trades transacted under relational contracts confirms and calibrates the hypotheses forwarded by both BFF and CR, i.e., welfare exceeding self-interested Nash equilibrium levels is driven, in part, by the presence of subjects with other-regarding preferences. Furthermore our analysis refines the contribution of BFF by identifying how specific types of social preferences shape the ebb and flow of contractual negotiations and performance.

We find a substantial role for social preferences in explaining the type of trade that is observed (public vs. private), the content of offers (price and effort requests), sellers' performance and buyers' expectations of seller performance, even after controlling for potentially confounding demographic determinants such as academic major and gender. In many instances the nature of social preferences exhibited during the CR games manifests in a straightforward manner in the contractual setting, e.g., social efficiency sellers improve total trading surplus by shirking less while buyers, who are vulnerable to shirking, fare relatively well in contracts when either party exhibits maximin social preferences.

Other manifestations of social preferences are more nuanced. For example, social efficiency and maximin buyers are much more likely to engage in private rather than public exchanges, while buyers displaying competitive preferences avoid private trades. This sorting of social preference types into different trading domains has fundamental implications for the generation and distribution of surplus because, on average, trades initiated via private offers feature more than twice the total surplus of publicly initiated trades and, furthermore, the surplus generated under private trading is more evenly divided between parties. The commitment of social efficiency buyers to utilizing private trades, coupled with infrequent shirking by social

efficiency sellers, results in frequent matching of social efficiency buyers with social efficiency sellers, which results in trades that often generate the maximum possible social surplus.

Many aspects of the resulting trade differ according to whether the parties are involved in privately or publicly initiated trades, with the role of social preferences being stronger in private trades. For example, social efficiency sellers shirk significantly less than other types in private trades, but shirk as much as other types in public trades. When trading is initiated privately, the buyer's type manifests most strongly in the non-binding quality request: once we control for relevant information concerning past performance of the seller and the existing length of the relationship, we find that competitive and inequality averse buyers set a high bar for seller performance while selfish buyers set the lowest bar. By contrast, in the public domain, quality requests differ by buyer type in the early rounds, but these differences disappear with time. Buyers instead begin to differentiate public offers via price, with selfish buyers setting the lowest prices holding past performance and other factors constant.

Our results also suggest a link between a subject's social preferences and beliefs. Specifically, after sellers have agreed to contract terms but before sellers actually deliver quality, each buyer is prompted to forecast the quality that will be delivered. Maximin buyers are significantly more optimistic about their seller's intentions than are other buyers both in public and private trades. The role of social preferences in forming these expectations persists, even after the level of promised performance and past transgressions against the buyer are controlled.

Our results also allow for decomposition of the efficiency wage effect. Like BFF we find that higher payment induces higher performance and that this holds even for subjects classified as purely selfish. After controlling for price, however, we find that the social preference type of the seller still explains a significant portion of performance. In public trades, for example, a buyer can offer a price to maximin seller that is 43 percent lower than that offered to a selfish

seller and expect the same level of performance. Maximin sellers' higher performance is consistent with helping the disadvantaged player as buyers regularly earn lower surplus in these contracts.

The results provide a detailed view of how heterogeneous social preferences may manifest in complex economic interactions. While much of the extant research implicitly acknowledges that heterogeneity exists with regard to the type and strength of social preferences that are present in experimental populations, much of this work focuses on identifying the social preference structure that best captures the average respondent. In contrast our approach is among the first to fully exploit subject heterogeneity by using a within-subject experimental design to identify how different social preferences manifest in alternative settings (later we discuss in detail the work by Mariana Blanco, Dirk Engelmann and Hans-Theo Normann (2006), which uses a similar approach). Finally, our research provides some insights into the link between *generalized* and *personalized* trust. Steven N. Durlauf and Marcel Fafchamps (2005) refer to personalized trust as trust that emerges from repeated games and is characterized by self-enforcing agreements and high levels of relationship specific surplus. In contrast, generalized trust refers to one's optimistic beliefs about a population of agents and is considered by Durlauf and Fafchamps to yield more efficient outcomes than personalized trust because generalized trust is established faster and more cheaply than personalized trust. Our results suggest that generalized trust is at least partly determined by social preferences as social efficiency types prefer to engage in private trades, shirk less than other types, and maximin buyers are significantly more optimistic than other buyers.

The remainder of the paper is organized as follows. In section 2 we discuss our experimental approach and design. In section 3 we provide a summary of results concerning the social preferences of our experimental subject pool, while in section 4 present results concerning

the role of these social preferences in explaining the relational contracting results. The final section discusses these results and concludes.

2. Experimental Design

To explicitly link individual social preferences to relational contracting results we wed two established experimental protocols via a within-subject design. First a protocol adapted from Charness and Rabin (2002) is used to classify each subject's social preferences. The adapted CR protocol pairs subjects together and requires 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). By having each subject participate in a sequence of games with variations in roles and possible outcomes, various social preferences can be identified. We consider six classes of social preferences: pure self interest – maximizing own payment; competitive – ensuring own payment is greater than the payment received by the other player (hereafter, other payment); negative reciprocity – choosing an option that reduces first-mover payment if the first mover's choice will reduce second-mover payment; disadvantaged inequality aversion – ensuring payments are as equal as possible when own payment is less than other payment; maximin – choosing the option where the least-well-off player's payment is highest; and social efficiency – maximizing the sum of own and others' payments. [1]

The CR 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

[1] Several authors also consider positive reciprocity, i.e., situations where a 'good' move by the first actor inspires the second actor to reduce own pay to help the first actor, and advantaged inequality aversion, i.e., ensuring payments are as equal as possible when own payment is greater than other payment. Empirically, these motivation have not been found to be as strong or prevalent (Charness and Rabin, 2002 and references therein). Hence, we omit then from consideration.

mover on subsequent occasions during any part of the experiment. CR recognize that '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 (Eldar Shafir and Amos Tversky (1992), Rachel Croson (1999); Werner Güth, Steffen Huck and Wieland Müller (2001); Stephen V. Burks, Jeffrey P. Carpenter and Eric Verhoogen (2003)) or similar (Timothy N. Cason and Vai-Lam Mui (1998), Brandts and Charness (2000), Robert J. Oxoby and Kendra N. McLeish (2004)) measures of underlying social preferences. Hence, our use of social preference estimates derived from the CR protocol can test the validity of using 'cold' social preference measures to predict behavior in 'hot' market settings.

The second protocol is a contractual economy adapted from Brown, Falk, and Fehr (2004). To facilitate inter-study comparison, the experimental market environment is programmed in Z-TREE (Fischbacher 1999) using base code employed by BFF.

In this experimental economy, trade of a vertically differentiated good among subjects is conducted via bilateral contracts. Subjects are randomly partitioned into two groups: buyers and sellers. Buyers offer contracts to sellers specifying a price-quality combination for a unit of an abstract good.[2] An offer can be made to all sellers simultaneously with public knowledge that all other sellers are receiving the exact same offer (public offer) or to a single seller with knowledge that no other seller is receiving this offer (private offer). A buyer can submit as many public and private offers as desired, though the first acceptance of any offer by a seller automatically withdraws all the buyer's remaining offers, i.e., the buyer may only contract with one seller per period. Sellers can accept, at most, one offer.

[2] Experimental directions use the terms buyer, seller, price and quality, though we will interchangeably discuss results using the terms employer, employee, wage and effort.

Sellers outnumber buyers, hence some sellers will not participate in each round and instead collect a reservation payment, i.e., there is involuntary unemployment. Furthermore, the seller can shirk, i.e., deliver quality lower than the agreed upon level, without contemporaneous financial penalty. In other words, seller performance is unenforceable. Buyers must pay the agreed-upon price, but need not make an offer to a seller in subsequent rounds.

Relevant market parameters mirror those used by BFF. Buyers' per-round payouts are:

$$(1) \qquad \pi_b = \begin{cases} 10Q - P & \text{if trade occurs} \\ 0 & \text{if no trade occurs} \end{cases}$$

where π_b is the buyer's payment, Q is the delivered quality level, and P is the agreed upon payment. Sellers' profit is:

$$(2) \qquad \pi_s = \begin{cases} P - c(Q) & \text{if trade occurs} \\ r & \text{if no trade occurs} \end{cases}$$

where r is reservation earnings and $c(Q)$ 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 quality (Q^a) and delivered quality (Q) are allowed to be any integer from one to ten while the price can be any integer from zero to one hundred.

The trading environment allows for reputation formation. Specifically, each party retains a unique identification number across all rounds and, at the end of each trading round, each subject is informed of the delivered quality and the payoffs for each partner (buyer and seller). 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 (i.e., develop personalized trust) over time while performance-contingent renewal can be wielded as a disciplinary 'stick' to foster that trust.

Each contracting session is comprised of two practice rounds and fifteen rounds that may determine eventual cash payment. As in BFF, practice rounds featured only the bidding stage of contracting to avoid costless, deceptive signaling by sellers.

Our protocol differs from BFF in several ways. First, due to limitations on computer lab size, our protocol features a thinner market: five buyers and seven sellers compared to seven buyers and ten sellers in BFF. The ratio of buyers to sellers is virtually identical, however. Second, our subjects participate in two contracting sessions per experiment rather than one, i.e., a relational contracting session and a session that is identical in all facets except quality is perfectly enforced.[3] The order of the two sessions within an experiment is counterbalanced across experiments to compensate for such differences; all identification numbers of subjects are changed between contracting sessions, and subjects are made aware of this; and experiment-specific dummy variables are used to control for order effects in regression analyses. Third, contractual rounds were shortened from 3 minutes to 90 seconds.[4] Fourth, the average level of compensation in our experiments is about $23, which, after adjusting for differences in local costs of living and differences in the length of the experiment, is slightly more than half the compensation rate received by average BFF subjects. Finally, to reduce income effects that might emerge during a second session, one experimental market session is chosen to be the 'paying' session via a publicly observed random process after both sessions are complete (a similar random incentive tactic is employed among the CR decisions rendered by subjects).

[3] These sessions mirror BFF's complete contracting sessions. Each session is identical to the relational contracting session except the computer enforces delivery of the agreed-upon quality, i.e., $Q^a = Q$.

[4] A pilot session was conducted that featured longer rounds. Nearly all trades were transacted within the first 90 seconds of the round. Furthermore, nearly all the 90-second sessions included in this study featured full trading. We conclude that the shorter rounds had little influence on resulting subject behavior and our results.

Subjects are fully informed of this compensation tactic at the beginning of the experiment and advised to treat each session with equal seriousness.

Eighty-four subjects participated in seven two-hour experiments that featured two contracting sessions and eight social preference games. Subjects were either a buyer or seller for all games in both contracting sessions, while each subject played both roles within each CR game. All trading takes place on networked computers enclosed in cubicles to eliminate between-subject visual contact. Social preference questions were divided in to two packets of four games with one packet administered prior to each contracting session; the order of the packets and the order of games within packets were counterbalanced across subjects and across experiments. Subjects were students recruited via e-mail and newspaper from various academic departments at the authors' home institution.

Our within-subject design is closest to that of Blanco, Engelmann and Normann (2006, hereafter BEN), though it differs in several ways. BEN's design requires subjects to cast seven decisions during four simple games (ultimatum, dictator, public goods and sequential prisoner's dilemma games), where second-mover's decisions cast in the two games featuring sequential play (ultimatum and prisoner's dilemma) were elicited via the strategy vector method. BEN use 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.

3. Social Preference Results

In the course of participating in the eight CR games described in Table 1, each subject cast thirteen decisions (three games were dictator games and featured a decision for only one role). Compared to responses from similar games reported in Charness and Rabin (2002), six of thirteen responses yield frequencies that are statistically similar to those of CR's subjects, who

were drawn from university communities in Berkeley, California, and Barcelona, Spain. For the remaining seven games, our subjects made decisions more consistent with pure self interest.

Decisions cast during each CR game can be interpreted as consistent with one or more classes of social preference described in the previous section. For the eight decisions in which the subject chose last or cast the only decision (e.g., dictator games), the consistent social preference classes are listed in the last column of Table 1. For games in which the subject was the first of two movers, we do not infer social preference classes because we lack subjects' beliefs regarding how the anonymously-matched partner would decide. Without these beliefs it becomes difficult to restrict the classes of social preferences consistent with a particular decision. Hence, for the remainder of the paper, we ignore first-mover responses to CR games.

We next classify each subject's 'closest' social preference class by comparing observed behavior and behavior predicted by the utility functions associated with the six social preference classes (see the appendix for more details and James Andreoni and Lise Vesterlund 2001 for a similar classification tactic). We find 45 percent of subjects perfectly adhere to the predictions of one of the six social preference classes (Andreoni and Vesterlund (2001) find this for 43 percent of their sample); the remaining subjects are assigned to the closest class.

Half of our subjects are classified as selfish, while about equal numbers are classified as competitive, disadvantaged inequality averse, social efficiency, and maximin (Table 2). Only a few are classified as reciprocal. Our proportion of self-interested types aligns with Andreoni, Marco Castillo and Ragan Petrie (2003) who found about half of their sample adhered to pure self interest in a variant of the ultimatum game.

Given the fact that some games are played after exposure to a contracting session, there was the possibility that the market experience influenced subsequent CR game decisions and our social preference measurement. Indeed, Brandts, Riedl and van Winden (2004) find evidence of

such feedback in an experimental setting where social preferences are measured before and after exposure to different degrees of market competition. However, for our subjects and this particular setting, we find no statistical evidence of such an influence for the eight CR responses used in this analysis. Specifically, multivariate analysis of variance tests for responses to the eight games confirm there is no difference in responses by the subject's role during the contracting game (buyer versus seller, $p = 0.29$) or by order of game administration (before or after the first contracting session during a given experiment, $p = 0.72$).

4. Relational Contracting Results

When quality is unenforceable, it is straightforward to show that an equilibrium involving purely self-interested subjects predicts full trade, minimal quality delivery ($Q = 1$) and low price ($P = 5$); hence, higher quality and price are suggestive of deviations from pure self interest. To see this, note that in a one-shot interaction, the sequence of events are that (1) the buyer offers a contract which specifies P and Q^a, and (2) if the seller accepts, he chooses Q, which can deviate from Q^a. Since the seller moves last and P is guaranteed, the seller maximizes that period's profits with $Q = 1$. The buyer, anticipating the seller's behavior, sets $P = 5$, which ensures that the seller earns his reservation utility. If instead, the parties interact repeatedly and it is common knowledge that the game will end after round fifteen, then the same outcome prevails in each round. To see this, note that, in the final round, the parties are essentially in a one-shot situation so the one-shot outcome prevails. By backward induction, the same outcomes occur in all previous rounds. The level of individual and joint surplus at the pure-self-interest Nash equilibrium is significantly lower than the social optimum of $Q = 10$.

BFF provide an analytic example where a subject pool consisting of 60 percent 'fair' subjects can sustain a socially efficient equilibrium ($Q = 10$) until the third to last period before deteriorating during the final periods. While their experimental results do not approach $Q = 10$,

they find quality increases from six in the first round to eight by the thirteenth round before declining thereafter (BFF's Figure 5). A similar time trend emerges in our data (Figure 1), though the absolute level of quality is lower than that observed by BFF.

4.1 The Type and Number of Buyers' Offers

The buyer sets the tone for subsequent interactions by proposing one or more offers where each offer is attached to her unique, round-invariant identification number. Analysis of buyers' offers provides an opportunity to isolate the genesis of social preferences' impact on market behavior. By the design of the computer interface, a buyer's first decision is whether to make an offer 'private' or 'public.' Both types are used extensively (45 percent public). Private offers enhance the ability to establish long-term relationships. For example, if a buyer wants to establish a relationship with seller A, then she only needs to make a private offer to seller A in every round rather than venture into the open market with a public offer and hope that seller A is the first to accept. This is an essential feature of relational contracting as it implies that the promise of future *relationship-specific* gains from trade can sustain high performance in the current period; i.e. it implies the existence of personalized trust.

Public offers, on the other hand, instantaneously communicate with all sellers, including sellers with whom the buyer has no past experience. In contrast a private offer is viewed only by the seller who is extended that offer. In both public and private offers, however, buyers only have knowledge of sellers with whom they have previously interacted, and have no way of identifying the past performance or activity of other sellers.

The use of public contracts is significantly different across the six classes of buyers (Table 3, row 1, $\chi^2(5) = 21.68$, $p<0.001$). Three of four offers issued by competitive and inequality averse buyers are in the public domain, while social efficiency and maximin buyers issue only 29 and 13 percent of offers in the public domain.

To control for other potential demographic and lagged behavioral influences upon the type of offer issued, an econometric model of the probability of issuing a public offer is presented in Table 4.[5] Selfish buyers serve as the base category for the social preference categorical variables. The regression controls for several key demographic characteristics including gender (69 percent male), age (21.2 years ± 3.9), race (69 percent white), academic major (44 percent are business and economics majors, 15 percent are other quantitative majors such as engineering and mathematics, and the remainder are some other major) and grade-point average (3.26 ± 0.51). Hence, our measurement of the influence of social preferences upon contracting behavior is not confounded with influences that demographics might have upon social preferences (e.g., gender, Andreoni and Vesterlund 2001; age, John A. List 2004; academic major, John R. Carter and Michael D. Irons, 1991). Furthermore, following BFF, we include controls for lagged factors that might influence the propensity to extend a public offer, including lagged seller performance and a time trend.

To examine the influence of social preferences on the probability of making a public offer, we test for the joint significance of buyer social preference parameters ($\chi^2(5) = 37.45$, $p < 0.01$). We also report the pseudo-R^2 of the model when buyer social preference class is dropped. The pseudo-R^2 drops to 0.31 from 0.41, which suggests a substantial role for intentional measures of social preferences in explaining the type of offer extended by buyers.

Buyers who display maximin and efficiency preferences during intentional CR games are significantly less likely to extend a public offer than selfish buyers (the omitted category) even after controlling for factors such as demographics and trading history that also influence the propensity to enter into private relationships. Figure 2 plots the percent of trades in each period

[5] All econometric models are presented with robust standard errors that are clustered on the 155 unique buyer-seller pairs that are observed during the 512 trades to account for unobserved effects from the pairing of two subjects, which might cause the composite error terms from the observations within each pairing to be correlated.

that are public for selfish buyers, for maximin and efficiency buyers combined and for competitive and inequality averse buyers combined. Selfish buyers rely more heavily upon public trades in nearly every period than do social efficiency and maximin buyers. Competitive and inequality averse buyers issue a similar percent of public offers as other types of buyers in the first round but show no trend toward private relationships over time like other types.

In addition to buyer social preferences, the tendency to extend a public offer hinges tightly on a buyer's past experience. For example, we find that the probability of extending a public offer increases when last period's seller failed to meet performance expectations[6] and decreases as the tenure of an existing buyer-seller relationship lengthens.

One possible implication of these results is that social preferences provide an exogenous explanation for when personalized trust is most likely to be established. Maximin types consistently seek to establish personalized trust, perhaps because they have more generalized trust than other types, and that selfish and particularly competitive types are much less likely to engage in long term relationships built on trust. Selfish and competitive buyers appear to prefer to operate through the public market place rather than form bilateral, self-enforcing agreements.

4.2 Buyer Price Offers and Quality Demands

Once the domain of the offer is established (public vs. private), a buyer specifies a quality request, which the seller need not follow, and a price, which must be paid by the buyer regardless of seller performance. The price is crucial to the seller because it sets the maximum possible payment obtainable by the seller (i.e., the seller can fully shirk at zero cost and collect full price).

Given the pivotal nature of the offer's content and the functional differences in public versus private offer mechanisms, it is not surprising that we find offers extended via public exchange are distinct from private offers. Public offers are lower in average price (22.73 vs.

[6] In each round buyers are asked to predict the performance of their seller after the seller has accepted an offer but prior to actual performance.

42.11, $p < 0.001$), average quality requested (6.18 vs. 8.57, $p < 0.001$) and the average percent of surplus designated for the seller (34.72 vs. 45.06, $p < 0.001$).[7]

Clearly the domain of contracting affects the content of the offer. However, the buyer's decision concerning the domain is likely to be endogenous. Hence, separate models are estimated for public and private offers as part of an endogenous two-stage switching model where the public offer model in Table 4 serves as the first stage and the inverse mills ratio from this model is included as a regressor in the Table 5 models.[8]

Within each domain of offers, social preferences play a role, though the role is substantially larger when private offers are considered. Specifically, the null hypothesis that buyer type does not affect price or quality request is rejected for private offers (Table 5, last row). Among public offers, however, it is rejected in the model of price only.

Among private offers the largest degree of differentiation across social preference types involves the level of quality demanded (the pseudo-R^2 declines more than 15 percent when buyer type is excluded). Recall the buyer's quality demand does not directly affect the size or distribution of surplus because sellers have unilateral latitude to shirk. However, quality requests may play an important role in setting expectations for renewal, i.e., mechanisms that sort and match buyers and sellers throughout the course of trading. For example, sellers who deliver a quality level of 4 to a buyer who demanded 4 may expect the buyer to renew next period while renewal may be less assured if buyer demanded more than 4.

Selfish buyers demand relatively little from their private partners – the average request is 8.24 quality units, which is significantly less than all other buyer types ($p = 0.03$). Competitive and inequality averse buyers demand substantially more. Because higher quality increases the

[7] Unless otherwise noted, all non-regression two-group comparisons reported in the text are two-sided Mann Whitney non-parametric tests.

[8] The exclusion restrictions used for the first-stage probit is the lagged 'negative surprise' variable, which takes on a value of one if the seller delivered quality lower than the buyer expected and zero otherwise. The mechanism for determining buyers' expectations is discussed in footnote 6.

payoffs of buyers relative to sellers, everything else equal, it makes intuitive sense that competitive and inequality averse buyers request more quality.

When raw averages across buyer preference types are computed (Table 3), social efficiency buyers ask private trading partners for higher quality than the average of other buyer types (8.82 vs. 8.48, $p = 0.08$) and selfish buyers in particular (8.82 vs. 8.24, $p = 0.04$), which makes sense given that social efficiency is a monotonic function of quality delivered and that, as we will see later, higher quality requests result in higher quality deliveries holding all else equal. In the regression results, however, social efficiency buyers' quality requests are not significantly higher than those of selfish buyers. However, in the regression model for quality request, lagged quality delivery has a significant, positive influence on requested quality. Social efficiency buyers sustain private relationships with higher lagged values for quality than other buyers (7.63 vs. 6.49, $p = 0.01$, not reported in Table 3), which likely drives the lack of significance for the social efficiency dummy variable in quality request regression model.

There is less statistically significant differentiation with respect to the prices extended in private trades – the regression model predicts that average buyers of most types will make offers in the low to mid-40's. There is a statistically significant trend toward maximin buyers' offering lower prices than selfish buyers, however. Sellers receive higher per-period payments on average than do buyers (26.0 vs. 17.9 for all trades, $p < 0.001$ and 31.1 vs. 27.7 for private trades, $p = 0.18$).[9] Hence, the stinginess of maximin buyers with regard to price does not directly diminish the welfare of the worse-off partner (i.e., themselves).

Another way to note the differences across social preference classes is to consider the distribution of the surplus that each type of buyer is proposing to her partner. Within private offers, a regression using the same explanatory variables as those in Table 5 (results are not

[9] The reported p values are for a Wilcoxon matched-pairs signed-rank test for the equivalence of buyer and seller profits.

presented) shows that competitive and maximin buyers offer significantly less generous offers than do selfish buyers in terms of the percent of total proposed surplus designated for the seller. This is fully consistent with competitive types attempting to receive a larger payment than their partner. With regard to maximin buyers, again, this is not inconsistent with attending to the needs of the worst-off player (themselves). Note that our result that selfish buyers make positive surplus offers to sellers appears to contradict an earlier prediction that selfish buyers will only offer a price sufficient to ensure that the seller earns her reservation utility. It should be borne in mind, however, that if selfish buyers believe that there exist non-self regarding types who are willing to cooperate even in the final period of our finitely repeated trading environment, then even selfish buyers may cooperate in earlier periods and make generous offers to sellers. Thus, our results do not contradict well known models of reputation building (e.g. the model of David M. Kreps, Paul Milgrom, John Roberts, and Robert Wilson (1982)) or previous results by BFF, which suggest that if rationality or selfishness of all traders is not common knowledge, then even selfish types can be disciplined in a finitely repeated game.

The regressions in Table 5 also clarify that private offers are strongly influenced by past experience with the seller, regardless of the social preference structure of the buyer. As the relationship lengthens, and as past quality performance improves, the buyer increases the requested quality and opens her purse strings. While price responds positively to the seller's past performance, it is not an overwhelming response. An additional unit of quality delivered by the seller in the previous period yields, on average, an additional 1.8 units of payment, while an additional 0.975 units of payment accompanies any renewal with the same buyer, yielding an increase of about 2.8 units in price (which is about seven percent of the private offer mean price of 42.1). Additional units of quality, however, cost the seller 1, 2 or 3 units of payment, with higher cost increments occurring for higher levels of absolute quality. Hence, at the margin, the

seller can only justify improving quality at low and middle levels of absolute quality. Indeed, less than five percent of sellers ever deliver quality above 8 units, the threshold at which the incremental cost of quality shifts from 2 to 3 units.

Once buyers venture into the public domain, however, the offers are more homogeneous across buyer types in general, though the general model masks some interesting learning trends across buyer groups. With regard to the quality request, the regression indicates no statistically significant difference across social preference types when all periods are considered. Additional regression results not reported here show that there is differentiation across buyer type during initial public trading – competitive and maximin buyers demand significantly lower quality during the first four rounds.

The prices attached to these public offers are significantly different across all trading periods. In contrast to the differences in quality requests, however, the distinction across buyer classes increases during later trading periods. Indeed, the F-statistic testing the null hypothesis that buyer type has no impact on price is 4.10 when only the first five periods are used. It grows to 9.14 when the last five periods are used instead. Across all rounds, selfish buyers offer lower prices, on average, than any other group when all other factors are held constant, though this difference is only statistically significant when compared to competitive and negative reciprocal buyers. During the last five rounds, however, social efficiency buyers also offer significantly higher prices than selfish buyers while competitive buyers offer significantly lower prices than selfish buyers.

Taken together, this suggests that the evolution of public trading patterns differ across social preference types. Buyers' quality requests quickly converge, and, instead, different types of buyers differentiate with regard to prices. Higher prices may be offered as a carrot by all buyers other than selfish and competitive types with the hopes of enticing fair sellers into

accepting their contracts. As BFF noted, fair sellers may reject unfair offers; thus, higher prices by our non-selfish/competitive buyers may be means by which these buyers might increase the odds that non-selfish sellers will accept their contracts.

4.3 Social Preferences and Buyers' Expectations of Seller Performance

An important intermediate output of the computerized trading environment is the buyer's expectations concerning seller performance. After a seller's acceptance of the offer but before his subsequent performance, the buyer is prompted to predict the seller's quality delivery. Given that shirking is the norm (55 percent of trades involve shirking), a buyer may not consider a small deviation in quality that unusual and, therefore, we focus on buyers' expectations of quality. BFF show that deviations from a buyer's quality expectations do play a pivotal role with fewer renewals occurring if sellers fail to meet these expectations (BFF's Table 3).

Our buyers expect sub-contractual performance in 45 percent of trades with an average expected performance that is 1.5 units below the quality specified in their offer. As was the case for a buyer's price and quality demands, a buyer's expectations of shirking differ by the domain of the offer: buyers expect more shirking in publicly initiated exchanges (1.87 vs. 1.20, p = 0.002). As before, we estimate a two-stage endogenous switching model to capture the buyer's expectations concerning shirking in each domain (Table 5, columns 3 and 6).

After accounting for other important factors that can shape performance expectations, e.g., the quality request, price, past seller performance and the length of buyer-seller relationship, a buyer's social preferences explains about ten percent of the variation in both private and public contracts. The strongest regularity in terms of buyer type is that maximin buyers show a consistent optimism in both domains: they expect less shirking than any other social preference group and significantly less shirking than selfish buyers. In the public domain selfish and social

efficiency buyers are among the most pessimistic while competitive types are the most pessimistic in private trades.

From the regression it is clear that buyers expect price to induce higher quality, i.e., buyers believe in efficiency wages or gift exchange. The regression displays a quadratic pattern, with the effect of price upon quality expectations reversing around 62 for private trades and around 72 for public trades (90 percent of trades feature prices less than 62). Expectations of shirking also depend upon the trading round in our fixed horizon market. In private exchanges there is no time trend, i.e., pessimism does not grow or shrink as the rounds progress. However, buyers are not naïve: their expectation of shirking more than doubles in the final period – a time when relationships are forced to end due to the fixed horizon of the market. On the other hand, for public exchanges the degree of pessimism is not affected by the arrival of the final period. Instead, there is a marginally significant time trend where the expected quality shortfall steadily rises from 0.7 in the first round to 3.0 in the final round.

This model also reveals how buyers perceived the relative effectiveness of their quality requests in each domain. Holding price constant, a one unit increase in the quality request in a private exchange increases expected shirking by about a half unit. Once the mode of exchange is a public offer, buyers expect shirking to increase nearly 0.7 units for every additional unit of quality requested, suggesting the buyers realized the relative futileness of quality requests within public offers. This is congruent with our previous finding that buyers' quality requests in public offers became less heterogeneous as time progressed.

4.4 Renewal of Private Contracts

The BFF trading environment is particularly useful for studying relational contracts because it provides a mechanism to nurture longer-term relationships via renewal of private offers. A majority of private trades (56 percent) occur between a buyer and a seller that are

continuing a past, private trading relationship. The non-renewal of this relationship provides the contractual stick that may be wielded if sellers exhibit opportunism in the previous period.

To see how a buyer's social preference may impact the use of this key instrument, we estimate two models (Table 6). The first model postulates that private contractual renewal is solely a function of the buyer's social preference type, other buyer traits, a time trend, and fixed session effects. The probit regression model depicts no significant role for social preferences in general ($\chi^2(4) = 5.13$, $p=0.27$), though social efficiency buyers are significantly more likely to renew private contracts than selfish buyers ($p=0.048$).

The second model, inspired by BFF's renewal model, postulates that, along with the buyer's social preference type and other characteristics, the buyer's history with the seller influences the probability of renewal with lagged quality delivery (both absolute quality delivered and quality surprises) and the sheer length of history between the partners driving the renewal decision. Once these relationship specific items are controlled, there is no significant differentiation of renewals across buyer types ($\chi^2(4) = 4.39$, $p=0.36$) and no pair-wise difference between selfish and social efficiency buyers ($p=0.205$). Most of the explanatory power is derived from absolute lagged quality delivery (high quality → renewal), relationship length (longer → renewal) and a time trend (later period → renewal). Contracts that result in a negative quality surprise are renewed no less frequently than those that meet buyer expectations once the absolute level of quality delivery is held constant. This is not surprising given that more than 40 percent of private quality deliveries fall short of expectations.

4.5 Quality Delivery and Sellers' Social Preferences

As the second-mover within the relational contracting setting, and as the only player with latitude to shirk, we expect sellers' social preferences to play a crucial role in quality delivery and, therefore, in generating total surplus. As with buyers, we expect seller performance will

significantly differ between public and private trades. Indeed, average delivered quality is substantially higher in private exchanges (6.99 vs. 2.86, $p < 0.001$). A likelihood ratio test rejects the null hypothesis that model parameters are identical for the models of private and public quality delivery ($\chi^2(25) = 88.26$, $p < 0.001$ for model 1, $\chi^2(23) = 174.00$, $p < 0.001$ for model 2). Unlike our analysis of buyers' offers, we assume the type of trade (public vs. private) is exogenous to the seller. We believe that this is a reasonable assumption given that decisions regarding the type of trade are made prior to decisions on quality delivery. Sellers are in the minority and often only receive a single offer or all offers of the same type (e.g., public), leaving them few chances to influence the domain of trade.

BFF found that sellers respond to higher prices with greater effort, supporting the efficiency wage hypothesis. In this spirit we estimate two models. Model 1 (Table 7) posits that quality delivery is a function of prices and seller type, while model 2 drops price as an explanatory variable because, in principle, a fixed price that does not vary with quality in each round should have no incentive effects.[10]

The results from model 1 (Table 7) reveal the expected positive coefficient on price in both private and public trades, i.e., regardless of social preference class and the private versus public nature of the trade, higher price induces higher performance over most of the possible price range. Holding price constant, however, we find large, statistically significant impacts of seller type on quality delivery in both public and private trades. For example, holding the price constant, the model suggests that maximin sellers consistently deliver higher quality in both private and public trades than do selfish sellers. Given the precarious position of the buyer, i.e., having limited recourse for disciplining shirking sellers, a delivery of high quality helps buyers who, across all treatments, average lower payments than sellers. Thus, it appears maximin

[10] Inter-period incentive effects would also be ruled out in a finitely repeated game if all subjects are purely self-interested.

sellers' behavior in this contractual setting is consistent with their behavior in the CR games where they focus on improving the payment of the person in the least advantageous position. In fact our maximin sellers tend to reverse the tables in private trades: buyers achieve higher profits (33.7 vs. 19.9, $p < 0.001$). In the eyes of the buyer, the impact of trading with a maximin seller is substantial. In private trading a buyer can offer a maximin seller a price that is 16 percent lower than that offered to a selfish buyer without receiving lower average quality; in public trades this grows to 43 percent.

Interestingly, social efficiency sellers deliver quality that is statistically indistinguishable from that of selfish sellers in the models that include price as a regressor (model 1). When interpreting results for model 1, however, the totality of the role that seller type plays is not clear because, as we know from Table 5, a seller's past quality delivery will influence a buyer's subsequent price offer in a recursive fashion. Thus, there may be some incentive effects from inter-period price variations that are functions of past performance. However, past performance is likely to be highly correlated with sellers' type so that social preferences may have both a direct effect on quality delivery and an indirect effect via price.

In model 2 we remove price as an explanatory variable to focus on the direct total influence of seller social preference type on delivered quality. We find that social efficiency types clearly emerge as the highest performing class of sellers in private trades. Quality delivery is precisely where such sellers should excel because total surplus is solely dependent upon the level of quality delivered to the buyer. In other words, if a subject cares about total surplus than delivering the highest possible quality is the only way to accomplish this. Indeed, social efficiency sellers deliver the highest possible quality in 74 percent of all private trades (compared to 35 percent of all other seller types).

The tendency for high quality delivery on the part of social efficiency sellers does not spill over to public trades, however, where social efficiency and selfish sellers deliver similar levels quality even when price is removed from the model (public model 2). The lower quality delivery by social efficiency sellers in public offers appears to be only a minor drag on overall surplus realization, however, because only one in five trades featuring a social efficiency seller starts with a public offer. This is significantly less than all other groups combined (20.9 percent vs. 47.3 percent, $p < 0.001$). It suggests that once social efficiency sellers enter private trades they then remain in private trading arrangements by delivering high quality, which leads to renewal and increasing prices. Indeed, once in a private trade, social efficiency sellers are more likely to be renewed than all other seller types (68.8 percent vs. 42.2 percent, Pearson $\chi^2(1) = 11.01$, $p<0.001$) because of high quality delivery. Finally, the evidence suggests that it is social efficiency *buyers* that are the chief beneficiaries of private relationships with these highly productive social efficiency *sellers*. Specifically, 53 percent of social efficiency sellers' deliveries in private trades go to a social efficiency buyer. Put another way, social efficiency buyers have 24.7 percent of their private trades fulfilled by a social efficiency seller, while other buyers have only 7.7 percent of trades fulfilled by social efficiency sellers (Pearson $\chi^2(1) = 14.62$, $p<0.001$). In other words, social efficiency types appear to 'find each other' through the buyers' heavy reliance on private offers and through sellers' systematic delivery of high quality.

4.6 Surplus Measures

Significant differences across social preference type consistently emerge for surplus measures obtained during private trades, though these differences are much weaker or non-existent for trades initiated in the public domain (Table 8). In particular, it is striking that there are no significant differences by seller type in the total surplus generated under public trades. This suggests that the type of market setting may have important implications for the

manifestation of social preferences with mechanisms that facilitate reputation formation exacerbating underlying heterogeneity that might emerge due to different social preferences.

Another gauge of the impact of social preferences comes from simple regressions of surplus levels upon buyer and seller social preference classes, buyer and seller demographic traits, controls for the period of trade and fixed session dummy variables. Separate regressions are run for public and private trades, yielding six separate models though, to conserve space, the results are not reported. Dropping the buyer and seller social preference classification yields a modest decline in R^2 for the public surplus regression models (6 percent for total and seller surplus, 11 percent for buyer surplus). In the private surplus models, however, removal of the social preference classification results in a major decline in explanatory power, with declines of 42, 34 and 29 percent for the R^2 in the social, buyer and seller surplus models. In other words, the well being of both partners is relatively unaffected by the partners' social preferences when trading is initiated with public offers. Once in a private trade, however, individual well being substantially varies with the social preferences of both parties.

5. Discussion and Conclusions

In this paper we use a within-subject design to link social preferences measured in a cold, intentional setting to subsequent behavior in a hot, dynamic relational contract setting. The explanatory role of social preferences in the present work is substantial with evidence that certain social preference types, including social efficiency and maximin types, sort themselves into private relationships that are able to sustain higher, more equitably distributed levels of surplus. Even within the self-selected private trading regime, the social preferences of both parties substantially influence the level and distribution of surplus with the parties' social preference type providing more than 40 percent of the explanatory power in regression models of the total surplus generated during such trades.

Though the subjects that display self-interest tend to generate less total surplus than subjects of some other social preference classes, these subjects do generate significantly more total surplus than the draconian predictions emerging from a Nash equilibrium involving subjects adhering to pure self interest. This is consistent with the model of reputation building by Kreps, Milgrom, Roberts, and Wilson (1982) and the claims of BFF that if rationality or selfishness of all traders is not common knowledge, then even selfish workers can be disciplined in a finitely repeated game. The presence of cooperative or other-regarding traders implies that rents might exist even in the final period so that even selfish workers have an incentive to cooperate in earlier periods to preserve the possibility of capturing some of this rent. Our analysis takes this argument one step further as we find that even beliefs about others' selfishness or willingness to cooperative might be a function of social preferences. We have evidence that links social preferences to one form of belief held by buyers – the belief that a partnered seller will not cooperate and deliver sub-contractual quality. While most buyers expect shirking, maximin buyers are more optimistic than self interested buyers, even after accounting for the contractual terms (price and quality request) and for sellers' past transgressions. Future experimental designs may be able to parse finer results concerning the beliefs of individual subjects, perhaps by using social preference games to sort subjects into markets consisting of participants that are more homogeneous with respect to social preference class and by making this sorting procedure known to all participants prior to the beginning of trading.

Much of the empirical social preferences literature recognizes that agents are heterogeneous in this regard. To our knowledge, other than the current work, there are few studies that attempt to link individual differences in social preferences to behavior in other settings. The study of Blanco, Engelmann and Normann (2006) comes closest to our present design and intent, though the authors' conclusion concerning the robustness of individual social

preferences across settings is more pessimistic than ours. Specifically they find little predictive power at the individual level though aggregate level results are more consistent with their measurements of social preferences. Their greater degree of pessimism may be driven by differences in experimental design and choice of functional form for social preferences. Specifically they use two data points generated by each subject during an ultimatum game to calibrate the two-parameter Fehr and Schmidt (1999) model of inequality aversion. This calibrated model is then used to derive sharp predictions for individual behavior in three subsequent games. In contrast we use eight separate data points for each subject and use this information to allocate each subject into one of six different preference categories. Our social preference classification is then used to look for broader regularities in a much larger, richer set of data generated by these individuals.

While we find broad consistencies between subjects' behavior in simple, cold laboratory games and in richer, dynamic laboratory games, the question of remains: Does the structure or strength of social preferences observed in the laboratory translate into natural market or institutional settings? Indeed, evidence compiled by List (2006) comparing aggregate behavior of subjects drawn from the same population in both laboratory and natural market settings casts doubt upon the proposition that social preferences displayed in the laboratory are successfully transplanted to competitive market settings. We view our within-subject design tactic as an integral step in further exploring this crucial question. By providing a more nuanced view of social preferences, it might facilitate the process of identifying the mechanisms that create disparities between field and laboratory results. For example, one might conjecture that field and laboratory settings may trigger different social responses; i.e., if in some field environments, the norms are such that people expect others to behave competitively and in a self-regarding manner, then subjects may exhibit *concern withdrawal* (Rabin and Charness 2002) where subjects

withdraw their willingness to behave in a non-self regarding way.[11] If such norms do not exist in the laboratory, then concern withdraw may never manifest itself in the laboratory and subjects may appear to be less self regarding and exhibit gift exchange tendencies.[12] In short, understanding manifestations of the spectrum of social preferences in different contexts may provide another avenue for researchers to identify laboratory biases and improve laboratory experimental methodology and design to make these experiments more relevant for real world applications.[13]

[11] One can easily imagine that even people who have a tendency to be other-regarding or kind may withdraw the willingness to be other self regarding when they walk onto a used car lot, where salespeople are often perceived to be extremely self interested and opportunistic toward customers.

[12] Note that in this example, in contrast to List's main point, it is not the attenuation of social preferences that potentially reduces gift exchanges; rather, in moving from a laboratory to a field setting, it is possible that one type of social preference (i.e. concern withdraw) negates another type of social preference (i.e. positive reciprocity).

[13] Our point here is similar to Charness, Guillaume R. Frechette, and Kagel's (2004) that experimental outcomes can be quite sensitive to seemingly innocuous changes in design/presentation/procedures, and/or cultural elements. In fact, these authors find that gift exchange can be significantly diminished by the mere inclusion of a payoff table. Hence, like List (2006), albeit for different reasons, these authors show that gift exchange is not as robust as previously thought.

Table 1

Percent of Subjects Choosing Left Response to Social Preference Questions

Game[a]	Options	Columbus (N = 84)	Berkeley[b] (N)	Barcelona[b] (N)	Preference[c]
1B[d]	(400, 400) v. (750, 375)	82	50** [e] (32)	52** (48)	L: M, C, S, D R:E
2A	(100, 1000) v. let B choose	31	50† (32)	--	--
2B[d]	(75, 125) v. (150, 125)	37	34 (32)	--	L: C, N, D R: M, E
3A	(700, 200) v. let B choose	79	56† (32)	--	--
3B[d]	(200, 700) v. (600, 600)	58	22** (32)	--	L: C, S R: M, E
4A	(375, 1000) v. let B choose	29	54** (35)	--	--
4B[d]	(400, 400) v. (350, 350)	95	89 (35)	--	L: M, S, E R: N
5B[d]	(300, 600) v. (700, 500)	77	--	67 (36)	L: C, S, D R: M, E
6A	(750, 0) v. let B choose	74	47** (36)	--	--
6B[d]	(400, 400) v. (750, 375)	81	61* (36)	--	L: M, C, S, D R: E
7A	(500, 500) v. let B choose	56	41 (32)	--	--
7B[d]	(800, 200) v. (0, 100)[f]	77	88 (32)	--	L: M, S, E R: C, N, D
8A[d]	(550,550,550)[g] v. (600,300, 900)	45	54 (24)	--	L: M, C, D R: S, E

Notes: (a) The letter 'A' ('B') refers to the first-mover (second-mover) role in the game. (b) From Charness and Rabin, Table I. (c) Abbreviations are L (R) – left (right) option chosen, M – maximin, C – competitive, S – pure self interest, D – disadvantaged inequality aversion, E – social efficiency, N – negative reciprocity. (d) Represents a game where the subject is the last mover or only mover; these games are used in subsequent analysis. (e) **, *, † denotes the value in this cell is significantly different at the one, five or ten percent level from the Columbus result via a chi-square test. (f) Berkeley version featured a payoff for B of 0 rather than 100 for role B. (g) Berkeley version of this game featured 575 instead of 550 for all three players.

Table 2
Distribution of Social Preference Types

Type	**% of Subjects**	**Associated Game Patterns[a]**
Selfish	50	LLLLLLLL (2)[b], LLLLLLLR (10), LLLLLRLL, LLLLRLLL (3), LLLLRLLR, LLRLLLLR, LRLLLLLL (4), LRLLLLLR (14), LRLLLRLL, LRLLLRLR, LRRLLLLR (4)
Competitive	10	LLLLLLRL[c] (3), LLLLLLRR (4), LRLLLLRR (2)
Disadvantaged Inequality Aversion	9	RRLRLLRL, LLLLLLRL[c] (3), LLRLLLRL, LLRLLLRR, LRRLLLRL, LRRLLLRR (2)
Social Efficiency	13	RLLRRLRL, LRRLLRLR, LRRLRRLL, RLRLLRLL, RRLLRLLL, RRRLLLLR, RRRLLRLR, RRRLRRLR (4)
Maximin	14	RRRLRRRR, LRRLLLLL (4), LRRLLRLL, LRRLRLLL (3), RLRRRRRL, RRRLRLLL, RRRLRLRL
Negative Reciprocal	4	RRRLRRLL, LLLRLLRL, LRRRLLRL, LRRRLRRL

Notes: (a) The eight games considered are those featuring the second-mover's response or only mover's response. (b) The number of subjects following this pattern is in parentheses. If the number is omitted, only one subject had this particular pattern of choices. (c) Subjects with this response pattern were classified as disadvantaged inequality aversion and competitive types, i.e., the indicator value for subjects with this pattern was set equal to ½.

Table 3
Means by Buyer's Social Preference Type

	All Types	Selfish	Competitive	D. Ineq. Averse	Social Efficiency	Maximin	Negative Reciprocity	Median Test by Type[a]
% Public Offers	45.12	41.81	74.14**	75.00**	29.13**	13.33**	41.38	21.68**
For Public Offers								
Quality Request	6.18	6.46	6.12	4.47**	7.27*	7.75	7.33	11.71*
Price	22.73	27.03**	28.95*	11.62**	15.27*	21.25	26.50	29.28**
Expected Quality Shortfall	1.87	2.31*	1.12**	0.93*	3.50**	0.00†	1.17	15.18**
% Surplus to Seller	34.72	38.08*	40.34	22.57	39.75**	14.79	27.08	20.09**
N	231	97	43	45	30	4	12	
For Private Offers								
Quality Request	8.57	8.24*	10.00**	10.00**	8.82†	7.23**	9.59	14.71*
Price	42.11	43.80	23.73**	59.00**	43.74	25.62**	48.24	57.64**
Expected Quality Shortfall	1.20	0.87	4.87**	0.00**	1.38	0.46	1.94	9.43†
% Surplus to Seller	45.06	44.22**	6.99**`	50.00*	61.77	25.87**	38.47	51.34**
% Renewals	55.98	46.83**	28.57*	100.00**	71.21**	54.17	53.33	8.59
N	281	135	15	15	73	26	17	

Notes: †,*, ** denotes differences between the mean for the social preference type in that column and the mean of all other social preference types at the ten, five and one percent levels as determined by a Mann-Whitney two-sided test. (a) Nonparametric test of the equality of medians across the six social preference groups ~ $\chi^2(5)$.

Table 4.
Determinants of the Type of Offer Extended by Buyers

	Public Offer
Buyer Social Preference Type	
Competitive	0.736†
	(0.390)
Disadvantaged Inequality Aversion	-1.711†
	(1.037)
Social Efficiency	-0.969**
	(0.331)
Maximin	-1.362*
	(0.615)
Negative Reciprocal	-1.030
	(0.793)
Other Buyer Traits	
Male	0.060
	(0.281)
Age	0.045
	(0.037)
GPA	-0.887**
	(0.302)
Nonwhite	0.223
	(0.262)
Business/Econ Major	0.524
	(0.321)
Other Quantitative Major	-0.821†
	(0.432)
Other Explanatory Variables	
Lagged Negative Quality Surprise	0.352*
	(0.171)
Length of Contractual Relationship	-0.652**
	(0.106)
Trading Period	-0.037†
	(0.021)
Constant	2.465*
	(1.077)
Controls for Fixed Session Effects	YES
N	476
Ln(Psuedo Likelihood)	-190.75
χ^2	150.44**
Pseudo-R^2	0.41
Pseudo-R^2 without types	0.31
χ^2 that types = 0	37.45**

Notes: Probit regressions with robust standard errors clustered on buyer-seller pairs.
†, *, ** denote statistical significance at the ten, five and one percent levels.

Table 5.
Determinants of Content of Buyers' Contractual Offers and Expectations of Shirking

	Private Offers			Public Offers		
	Quality Request[a]	Price[b]	Expected Quality Shortfall[b]	Quality Request[a]	Price[b]	Expected Quality Shortfall[b]
Competitive	21.898**	1.472	2.435†	1.211	9.788*	-2.231**
	(3.922)	(12.024)	(1.255)	(0.896)	(3.952)	(0.534)
D. Ineq. Aversion	19.130**	22.720	1.040	-1.225	18.755	-0.104
	(2.913)	(12.496)	(1.019)	(2.316)	(11.748)	(1.485)
Social Efficiency	1.545	-0.341	0.607	0.529	15.198	0.992
	(1.249)	(6.918)	(0.439)	(1.761)	(11.644)	(0.815)
Maximin	2.308	-12.431†	-2.373**	3.795	12.044	-2.377*
	(2.146)	(7.403)	(0.839)	(2.420)	(7.546)	(1.168)
Neg. Reciprocal	5.667**	4.374	1.071	-1.413	17.197*	-0.590
	(2.100)	(11.761)	(1.011)	(1.798)	(7.456)	(0.903)
Male	3.138**	1.474	0.031	1.040	-4.975	0.553
	(0.979)	(4.056)	(0.391)	(1.005)	(4.539)	(0.623)
Age	-0.062	-0.011	-0.034	-0.153	-0.100	0.202*
	(0.173)	(0.617)	(0.050)	(0.138)	(0.770)	(0.081)
GPA	2.621*	9.389	0.421	2.386*	-1.643	-0.815
	(1.256)	(7.436)	(0.504)	(0.970)	(5.021)	(0.508)
Nonwhite	1.447	-8.434†	0.889†	1.604†	-8.889*	-0.238
	(0.962)	(4.471)	(0.453)	(0.884)	(3.986)	(0.520)
Business/Econ	-0.806	-0.086	0.925	-0.724	5.045	1.757**
	(1.058)	(7.116)	(0.568)	(0.886)	(5.594)	(0.654)
Other Quantitative	-0.209	10.844	0.447	-2.956	18.547	2.091*
	(1.411)	(7.668)	(0.505)	(1.850)	(11.317)	(0.939)
Lag Quality	0.609**	1.804*	-0.135**	-0.036	0.244	-0.067
	(0.114)	(0.731)	(0.045)	(0.116)	(0.523)	(0.086)
Rel. Length	0.570**	0.975*	-0.157*	0.102	1.725	0.157
	(0.180)	(0.484)	(0.062)	(0.399)	(1.135)	(0.157)
Price	--	--	-0.172**	--	--	-0.121**
			(0.030)			(0.044)

	Private Offers			Public Offers		
	Quality Request[a]	Price[b]	Expected Quality Shortfall[b]	Quality Request[a]	Price[b]	Expected Quality Shortfall[b]
$Price^2$	--	--	0.001**	--	--	0.001†
			(0.0002)			(0.0005)
Quality Request	--	--	0.468**	--	--	0.682**
			(0.084)			(0.078)
Trading Period	-0.185*	-0.673	0.025	-0.238*	-0.541*	0.110†
	(0.086)	(0.461)	(0.036)	(0.093)	(0.257)	(0.059)
Last Period	1.439	-2.615	1.157†	0.587	-4.487	-0.196
	(0.909)	(5.017)	(0.620)	(1.232)	(3.576)	(0.563)
Pub Offer Mills[c]	-0.033	0.362*	0.037*	0.159	-0.847	-0.063
	(0.038)	(0.155)	(0.016)	(0.168)	(0.768)	(0.088)
Constant	-4.650	9.157	2.082	3.340	54.535*	-2.127
	(3.646)	(16.796)	(1.628)	(4.803)	(26.680)	(3.351)
Controls for Fixed Session Effects	YES	YES	YES	YES	YES	YES
N	272	272	272	204	204	204
Ln(Psuedo Likelihood)	-314.44			-415.64		
χ^2	234.44**			116.94**		
R^2 or Pseudo-R^2	0.58	0.53	0.61	0.44	0.66	0.60
- without types[d]	0.49	0.50	0.52	0.42	0.62	0.53
χ^2 or F: types = 0[e]	92.47**	6.19**	9.05**	5.06	5.26**	7.24**

Notes: (a) Censored regression with robust standard errors clustered on buyer-seller pair and fixed session effect controls. (b) Linear regression with robust clustered on buyer-seller pair and fixed session effect controls. (c) Inverse Mills Ratio from probit regression of Public Offer. (d) The R^2 or pseudo-R^2 value that results from dropping all explanatory variables related to the social preference variables. The pseudo-R^2 for the zero-truncated negative binomial model is the R^2 between predicted and actual observations. (e) Chi-squared (for quality request model) or F-test (price and expected quality shortfall model) statistic associated with null hypothesis that all social preference type variables are jointly equal 0. †, *, ** denotes statistical significance at the ten, five and one percent levels.

Table 6.
Determinants of Buyer Decision to Renew Private Contracts

	Model 1	Model 2
Buyer Type		
Competitive	-0.854	-1.109
	(1.354)	(1.234)
Disadvantaged Inequality Aversion	--[a]	--[a]
Social Efficiency	0.789*	0.475
	(0.399)	(0.375)
Maximin	0.008	-0.264
	(1.003)	(0.814)
Negative Reciprocal	-0.076	-0.468
	(0.757)	(0.630)
Buyer Traits		
Male	0.652†	0.358
	(0.389)	(0.328)
Age	0.139†	0.134*
	(0.071)	(0.062)
GPA	-0.157	-0.410
	(0.433)	(0.433)
Nonwhite	-0.477	-0.294
	(0.428)	(0.366)
Business/Econ Major	0.086	0.042
	(0.459)	(0.387)
Other Quantitative Major	-0.117	-0.225
	(0.587)	(0.462)
Other Explanatory Variables		
Negative Quality Surprise	--	0.157
		(0.339)
Positive Quality Surprise	--	-0.114
		(0.456)
Quality Delivered	--	0.110*
		(0.046)
Relationship Length	--	0.146*
		(0.068)
Trading Period	0.414**	0.322**
	(0.118)	(0.117)
$(\text{Trading Period})^2$	-0.019**	-0.017*
	(0.007)	(0.007)
Constant	-5.103**	-4.234**
	(1.862)	(1.517)
Controls for Fixed Session Effects	YES	YES
N	237	237
Ln(Psuedo Likelihood)	-132.67	-122.04
χ^2	43.92**	175.35

	Model 1	Model 2
Pseudo-R^2	0.19	0.25
χ^2 that types = 0	5.13	4.39

Notes: Probit regression with robust errors clustered on buyer-seller pairs. †, *, ** denote statistical significance at the ten, five and one percent levels. (a) Disadvantage inequality aversion is dropped because it is perfectly (negatively) collinear with negative quality surprise.

Table 7
Determinants of Sellers' Quality Delivery

	---------- Private ----------		----------Public ----------	
	Model 1	Model 2	Model 1	Model 2
Competitive	1.946	-3.385	3.104*	3.367*
	(1.450)	(2.336)	(1.523)	(1.713)
Dis. Inequality Aversion	-1.978*	-2.415	-0.439	-0.480
	(0.847)	(1.751)	(1.142)	(1.196)
Social Efficiency	0.896	4.772**	0.718	1.416
	(0.890)	(1.586)	(1.279)	(1.631)
Maximin	1.279*	0.086	2.132*	2.230*
	(0.509)	(1.162)	(0.938)	(1.005)
Negative Reciprocal	-19.922**	-26.799**	-16.310**	-18.619**
	(2.012)	(3.070)	(1.965)	(2.121)
Male	0.004	-0.863	-1.097	-1.355
	(0.460)	(0.941)	(0.841)	(0.893)
Age	0.177†	0.420	0.193*	0.130
	(0.106)	(0.272)	(0.087)	(0.107)
GPA	-0.400	-1.217	0.047	0.688
	(0.562)	(0.970)	(0.821)	(0.907)
Nonwhite	-0.003	-1.110	1.536**	1.696**
	(0.542)	(1.381)	(0.469)	(0.620)
Business/Econ	-1.451**	-1.498	-2.091**	-2.375**
	(0.478)	(1.174)	(0.671)	(0.765)
Other Quantitative	0.200	-1.229	0.043	-0.485
	(0.664)	(1.234)	(1.094)	(1.238)
Price	0.259**	--	0.294**	--
	(0.067)		(0.047)	
$Price^2$	-0.001	--	-0.002**	--
	(0.001)		(0.047)	
Trading Period	0.040	-0.016	-0.003	-0.231**
	(0.056)	(0.091)	(0.060)	(0.070)
Last Period Dummy	-3.178*	-4.081**	-0.970	-1.308
	(1.251)	(1.416)	(0.826)	(1.388)
Relationship Length	0.784*	1.720**	-0.367	-0.619
	(0.366)	(0.559)	(0.758)	(0.937)
$(Relationship\ Length)^2$	-0.047	-0.077	0.185	0.296
	(0.035)	(0.051)	(0.295)	(0.423)
Constant	-5.476†	1.224	-6.901†	1.928
	(3.031)	(5.371)	(3.981)	(4.329)
Controls for Fixed Session Effects	YES	YES	YES	YES
N	281	281	231	231
Ln(Psuedo Likelihood)	-399.00	-480.96	-299.08	-326.24
χ^2	621.67**	323.33**	268.10**	225.03**
Pseudo-R^2	0.71	0.48	0.59	0.48
Pseudo-R^2 without types	0.67	0.40	0.55	0.44
χ^2 for types = 0	125.57**	110.35**	91.06**	93.31**

Notes: Censored regression with robust standard errors clustered on buyer-seller pairs and fixed session effect controls. †, *, ** denote statistical significance at the ten, five and one percent levels.

Table 8
Mean Surplus by Social Preference Type and Offer Type

	All Types	Selfish	Competitive	D. Ineq. Averse	Social Efficiency	Maximin	Negative Reciprocity	Median Test by Type[a]
By Buyer Type								
All Offers								
Total	43.9	46.2†	34.5**	34.7**	49.6*	41.6	44.1	17.3**
Buyer	17.9	17.3	11.7**	16.9	23.4*	22.3	12.5	9.5†
Seller	26.0	28.9**	22.8**	17.8**	26.2	19.3	31.6†	35.5**
% Public	45.0	41.8	74.1**	75.0**	29.2**	13.3**	41.4	65.4**
Public Offers								
Total	25.7	28.0†	33.8	20.5†	16.6**	25.0	20.1	12.8*
Buyer	5.9	4.3	9.6*	10.8	2.4	6.2	-4.8	9.3†
Seller	19.8	23.7**	24.2**	9.7**	14.2*	18.8	24.9	25.0**
Private Offers								
Total	58.8	59.4	36.7**	77.2**	63.2*	44.1**	61.1	52.4**
Buyer	27.7	26.7	17.6*	35.0**	32.0*	24.8	24.7	23.6**
Seller	31.1	32.7†	19.1**	42.2**	31.2	19.3**	36.4	54.4**
By Seller Type								
All Offers								
Total	43.9	43.6	28.2**	41.5	65.0**	42.3	10.0**	28.9**
Buyer	17.9	16.4†	14.3	6.4**	27.6**	25.1**	-32.8**	30.6**
Seller	26.0	27.2†	13.9**	35.1**	37.4**	17.2**	42.8*	83.7**
% Public	45.0	48.1	57.7	43.6	20.9**	43.6	60.0	13.4*
Public Offers								
Total	25.7	25.6	18.4	27.4	30.2	27.7	10.0	6.0
Buyer	5.9	5.2	3.7	-3.0	7.3	14.0**	-34.7**	8.5
Seller	19.8	20.4	14.7*	30.4*	22.9	13.7*	44.7*	10.5†
Private Offers								
Total	58.8	60.3	41.6*	52.3†	74.3**	53.6*	10.0*	48.5**
Buyer	27.7	26.8	28.7	13.6**	33.0*	33.7*	-30.0*	19.5**
Seller	31.1	33.5**	12.9**	38.7*	41.3**	19.9**	40.0	69.7**

Notes: †,*, ** denotes differences between the surplus measure for the social preference type in that column and the surplus obtained by all other social preference types at the ten, five and one percent levels as determined by a Mann-Whitney two-sided test. (a) Nonparametric test of the equality of medians across the six social preference groups ~ $\chi^2(5)$.

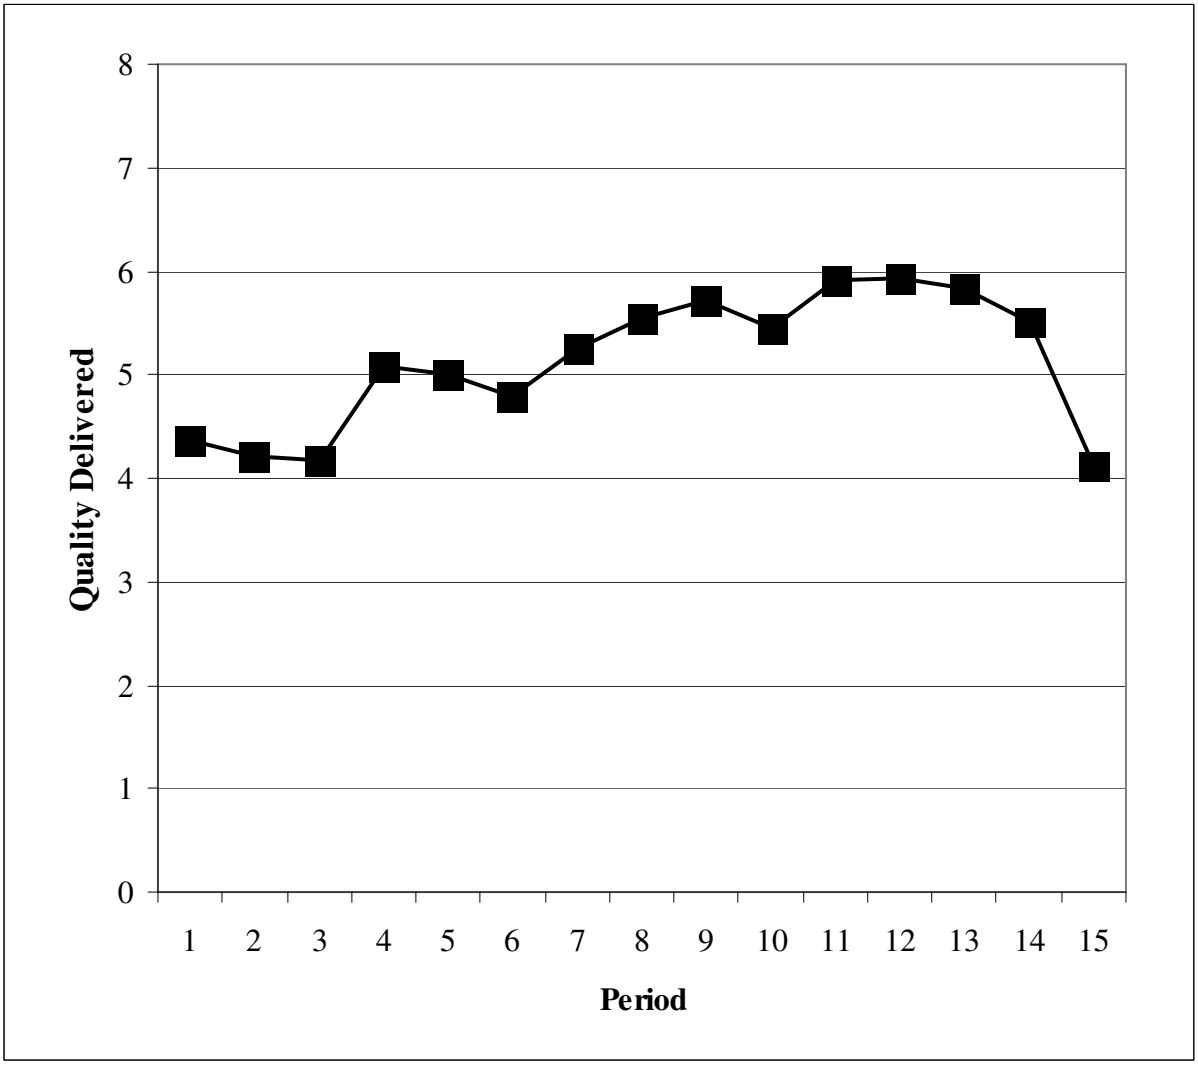

Figure 1

Average Delivered Quality by Round

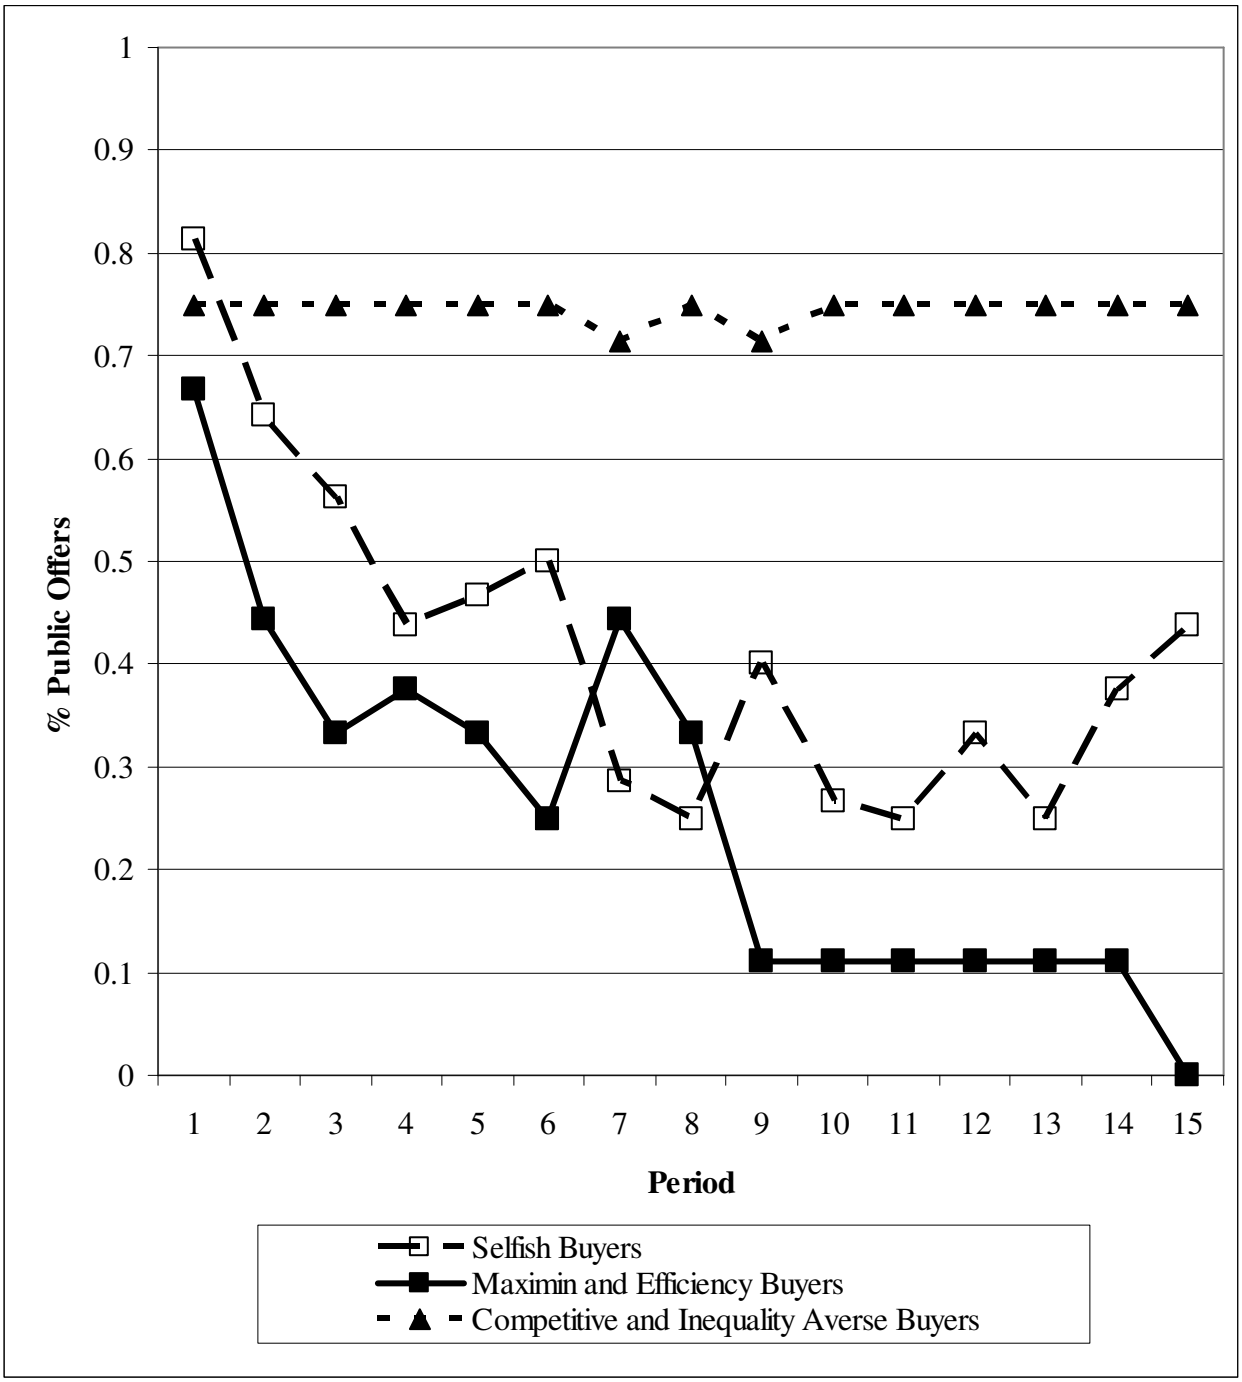

Figure 2.

Public Offers Made by Buyer Social Preference

References

Andreoni, James; Castillo, Marco and Petrie, Raga. "What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game." *American Economic Review*, June 2003, *93*(3), pp. 672-685.

Andreoni, James and Vesterlund, Lise. "Which is the Fair Sex? Gender Differences in Altruism." *Quarterly Journal of Economics*, February 2001, *116*(1), pp. 293-312.

Bereby-Meyer, Yoella and Niederle, Muriel. "Fairness in Bargaining." *Journal of Economic Behavior & Organization*, February 2005, *56*(2), pp. 173-86.

Blanco, Mariana; Engelmann, Dirk and Normann, Hans-Theo. "A Within-Subject Analysis of Other-Regarding Preferences." Working paper, Royal Holloway, University of London, 2006.

Bolton, Gary E.; Brandts, Jordi and Ockenfels, Axel. "Fair Procedures: Evidence from Games Involving Lotteries." *Economic Journal*, October 2005, *115*(506), pp. 1054-1076.

Brandts, Jordi and Charness, Gary. "Hot versus Cold: Sequential Responses and Preference Stability in Experimental Games." *Experimental Economics*, March 2000, *2*, pp. 227-38.

Brandts, Jordi and Charness, Gary. "Do Labour Market Conditions Affect Gift Exchange? Some Experimental Evidence." *Economic Journal*, July 2004, *114*(497), pp. 684-708.

Brandts, Jordi; Riedl, Arno, and van Winden, Frans. "Competition and Well-Being." Tinbergen Institute Discussion Paper, TI 2004-041/1, Tinbergen Institute, Rotterdam, Amsterdam April 2004.

Brown, Martin; Falk, Armin, and Fehr, Ernst. "Relational Contracts and the Nature of Market Interactions." *Econometrica*, May 2004, *72*(3), pp. 747-80.

Burks, Stephen V.; Carpenter, Jeffrey P., and Verhoogen, Eric. "Playing Both Roles in the Trust Game." *Journal of Economic Organization and Behavior*, June 2003, *51*(2), pp. 195-216.

Cason, Timothy N. and Mui, Vai-Lan. "Social Influence in the Sequential Dictator Game." *Journal of Mathematical Psychology*, June/September 1998, *42*(2/3), pp. 248-65.

Carter, John R. and Irons, Michael D. "Are Economists Different, and If So, Why?" *Journal of Economic Perspectives*, Spring 1991, *5*(2), pp. 171-177.

Charness, Gary. "Attribution and Reciprocity in an Experimental Labor Market." *Journal of Labor Economics*, July 2004, *22*(3), pp. 665-88.

Charness, Gary and Haruvy, Ernan. "Altruism, Equity, and Reciprocity in a Gift-exchange Experiment: An Encompassing Approach." *Games and Economic Behavior*, August 2002, *40*(2), pp. 203-31.

Charness, Gary; Frechette, Guillaume R. and Kagel, John H. "How Robust is Laboratory Gift Exchange?" *Experimental Economics,* June 2004, 7(2), pp. 189-205.

Charness, Gary and Rabin, Matthew. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics*, August 2002, *117*(3), pp. 817-69.

Cox, James C. "How to Identify Trust and Reciprocity." *Games and Economic Behavior*, February 2004, *46*(2), pp. 260-81.

Croson, Rachel. "The Disjunction Effect and Reason-based Choice in Games." *Organizational Behavior and Human Decision Processes*, November 1999, *80*(2), pp. 118-33.

Dufwenberg, Martin and Kirchsteiger, Georg. "A Theory of Sequential Reciprocity." *Games and Economic Behavior*, May 2004, *47*(2), pp. 268-98.

Durlauf, Steven N. and Fafchamps, Marcel. "Social Capital" in *Handbook for Economic Growth.* Phillipe Aghion and Steven Durlauf, eds., pp. 1639-1699, North-Holland 2005.

Engelmann, Dirk and Strobel, Martin. "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments." *American Economic Review*, September 2004, *94*(4), pp. 857-69.

Falk, Armin and Fischbacher, Urs. "A Theory of Reciprocity." *Games and Economic Behavior*, February 2006, *54*(4), pp. 293-315.

Fehr, Ernst and Falk, Armin. "Wage Rigidity in a Competitive Incomplete Contract Market." *Journal of Political Economy*. February 1999, *107*(1), pp. 106-34.

Fehr, Ernst and Fischbacher, Urs. "Why Social Preferences Matter – The Impact of Non-Selfish Motives on Competition, Cooperation and Incentives." *Economic Journal*, March 2002, *112*(478), pp. C1 – C33.

Fehr, Ernst; Gachter, Simon and Kirchsteiger, Georg. "Reciprocity as a Contract Enforcement Device." *Econometrica*, July 1997, *65*(4), pp. 833-60.

Fehr, Ernst; Kirchler, Erich, Weichbold, Andreas, and Gachter, Simon. "When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets." *Journal of Labor Economics*, April 1998, *16*(2), pp. 324-51.

Fehr, Ernst; Kirchsteiger, Georg, and Riedl, Arno. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *Quarterly Journal of Economics*, May 1993, *108*(2), pp. 437-60.

Fehr, Ernst; Kirchsteiger, Georg, and Riedl, Arno. "Involuntary Unemployment and Non-compensating Wage Differentials in an Experimental Labour Market." *Economic Journal*, January 1996, *106*(434), pp. 106-21.

Fehr, Ernst; Kirchsteiger, Georg, and Riedl, Arno. "Gift Exchange and Reciprocity in Competitive Experimental Markets." *European Economic Review*, January 1998, *42*(1), pp. 1 – 34.

Fehr, Ernst and Schmidt, Klaus M. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics*, August 1999, *114*, pp. 817-868.

Fehr, Ernst and Schmidt, Klaus M. "Fairness, Incentives and Contractual Choices." *European Economic Review*. May 2000, *44*(4-6), pp. 1057-68.

Fehr, Ernst and Schmidt, Klaus M. "Fairness and Incentives in a Multi-task Principal-agent Model." *Scandinavian Journal of Economics*, 2004, *106*(3), pp. 453-74.

Frey, Bruno S. and Stutzer, Alois. "Beyond Outcomes: Measuring Procedural Utility." *Oxford Economic Papers*, January 2005, *51*(1), pp. 90-111.

Güth, Werner; Huck, Steffen and Müller, Wieland. "The Relevance of Equal Splits in Ultimatum Games." *Games and Economic Behavior*, October 2001, *37*(1), pp. 161-69.

Hannan, R. Lynn; Kagel, John H. and Moser, Donald V. "Partial Gift-exchange in an Experimental Labor Market: Impact of Subject Population Differences, Productivity Differences, and Effort Requests on Behavior." *Journal of Labor Economics*, October 2002, *20*(4), pp. 923-51.

Kreps, David M.; Milgrom, Paul; Roberts, John and Wilson, Robert. "Rational Cooperation in the Finitely Repeated Presoners' Dilemma." *Journal of Economic Theory*, 1982, 27(2), pp. 245-252.

List, John A. "Young, Selfish and Male: Field Evidence of Social Preferences," *Economic Journal*, January 2004, *114*(492), pp. 121-149.

List, John A. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy*, February 2006, *114*(1), pp. 1 – 37.

Oxoby, Robert J. and McLeish, Kendra N. "Sequential Decision and Strategy Vector Methods in Ultimatum Bargaining: Evidence of the Strength of Other-Regarding Behavior." *Economics Letters*, September 2004, *84*(3), pp. 399-405.

Shafir, Eldar and Tversky, Amos. "Thinking through Uncertainty: Nonconsequentialist Reasoning and Choice." *Cognitive Psychology*, October 1992, *24*(4), pp. 449-74.

Appendix: Classifying Subjects into Social Preference Categories

Let $U^{i}_{j,g}(\pi_{j,g}, \pi'_{j,g})$ 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 six possible utility forms:

(Selfish - S) $U^{S}_{j,g} = \pi_{j,g}$,

(Competitive – C) $U^{C}_{j,g} = \pi_{j,g} - \pi'_{j,g}$,

(Disadvantaged Inequality Aversion – D) $U^{D}_{j,g} = -\max[\pi'_{j,g} - \pi_{j,g}, 0]$,

(Social Efficiency – E) $U^{E}_{j,g} = \pi_{j,g} + \pi'_{j,g}$,

(Maximin – M) $U^{M}_{j,g} = \min[\pi_{j,g}, \pi'_{j,g}]$ and

(Negative Reciprocity – N) $U^{N}_{j,g} = \psi$,

where ψ is a dummy variable that equals one if the first-mover implements an action that diminishes the subject's payment relative to other possible first-mover options and the subject then chooses an option that diminishes the first mover's payment (the three games in which this occurs are denoted in table I).

During the CR 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 ($\overline{V}^{i}$) 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 CR games, i.e., V^{i} =

$\sum_{g=1}^{8} U^{i}_{j,g}$. Normalized utility scores of the form $\nu^i = [V^i - \underline{V^i}]/[\overline{V}^i - \underline{V^i}]$ are then calculated for each subject and $\forall\ i \in \{S, C, D, E, M, N\}$. The subject is then assigned to the preference category that generates the highest normalized utility score.

Appendix: Experimental Instructions

Charness-Rabin Games

In this part of the session you will participate in several games. Each game is explained on a separate sheet that you will receive. You complete one game at a time until all the games are complete. Later in the session, you will participate in several more games like these.

In each game you have at most a single decision to make. In each game your decision, coupled with the decision of another person in the room, will affect your take home pay. You will never know the identity of the other person in the room who will affect the outcome of your game and you will never be matched with the same person twice.

Each game has 2 players – Player A and Player B. In each game there are several different combinations of possible payouts for Players A and B. In most games, Player A can choose one combination of payouts or pass up this option and let Player B choose from a separate menu of 2 different payout combinations.

In this part of the session there are four different games and each of you will play each game twice – once as Player A and once as Player B. If you are Player A, you can choose a certain outcome, or decide to let Player B determine the outcome. If you are Player B, you choose 1 of the 2 possible outcomes in case Player A let's you choose the outcome. Both players will be fully aware of each other's options in all games. Remember, for each decision you will be matched with a different person in the room and you will never know the identity of that person. Furthermore, you will not be informed of the results of any previous round or game prior to making your decision.

At the end of the night, you will be given the results of all decisions from this set of games and the similar set of games you will play later. Although you will have many 'outcomes' of the games, only one will be selected for payout. We will roll a die to determine which outcome will be added to your take-home pay. Because any of the outcomes can selected for payment, you should treat each decision as if it were the one that will be paid.

To convert the numbers on the decision sheets to dollars, use the following simple formula: 100 = $1.00.

At this point feel free to ask any questions. If there are no questions go ahead and start working through the packet.

Example of Charness-Rabin Game

ID#: _____ ROUND 1

GAME 2

In this ROUND, you are player B.

- You may choose B1 or B2.
- Player A has already made a choice.
  - If he or she has chosen A1, he or she receives 100 and you receive 1,000.
  - Your decision only affects the outcome if player A has chosen A2. Thus, you should choose B1 or B2 on the assumption that player A has chosen A2 over A1.
    - If player A has chosen A2 and you choose B1,

      → player A receives 75 and you receive 125.
    - If player A has chosen A2 and you choose B2,

      → player A receives 150 and you receive 125.

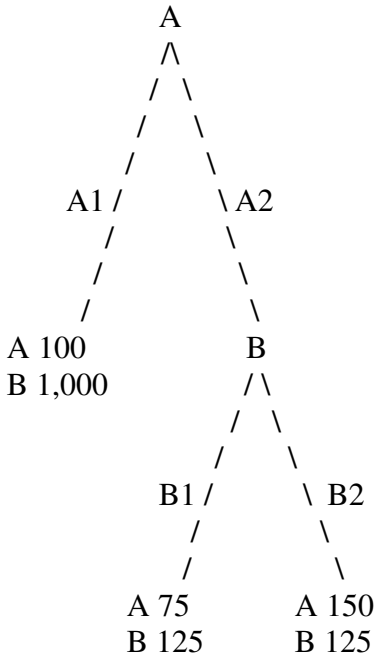

DECISION

I choose (circle one): B1 B2

Brown, Falk and Fehr Trading Game

Instructions for Buyers (Sellers' Instructions are similar and are not included)

In this experiment everyone begins with $5. During the experiment you can earn more money, with the exact amount depending on you and your pair member's decisions. During the experiment, your income is calculated in points. At the end of the experiment, points are converted into dollars at the rate of:

$1 = 70 points

Your initial balance of $5 equals 350 points.

Without exception, all written information you received from us is for your private use only. You are not allowed to pass over any information to other participants in the experiment. Talking during the experiment is not permitted. Violations of these rules would force us to stop the experiment. If you have any questions, please ask us.

General Information

The experiment is divided into periods. In each period, you have to make decisions, which you will enter on a computer screen. There are 15 identical periods in all and **the experiments ends at the end of period 15.**

Participants are divided into two groups consisting of 5 buyers and 7 sellers. **You will remain a buyer throughout the experiment.** This session will involve trading between buyers and sellers. The price agreed upon between the buyer and sell will determine how much money each party makes during the period.

Trades will take place on the computer screen. Buyers and sellers will each be identified by a number (from 1 to 7, e.g., buyer #3, seller #5, etc). Buyers and sellers keep the same identification number for all rounds of this game.

The Experimental Procedures in Detail

Each period is divided into a **trading phase** followed by a **quality determination phase.**

1. The Trading Phase

Each period starts with a trading phase. During the trading phase, each buyer can conclude a trade with one seller. In order to do so **each buyer can submit as many offers as he/she wishes.** In each trading phase, you will see a screen with the following features:

- The trading period is indicated at the top of the screen. The remaining time in the trading phase is also indicated at the top right corner. The **trading phase will last 90 seconds.** When the time is up, the trading phase is over and no further offers can be submitted or accepted.
- Once the above screen is displayed, the trading phase starts. As a buyer, you now have the opportunity to submit offers to sellers. Offers must include the following, which is to be entered into the right hand side of the screen:
  a) Specify whether the offer is to be public or private.

     Public offers will be communicated to all participants, both sellers and buyers. In turn, you will see all public offers by other buyers. **A public offer can be accepted by any seller. Simply click on the "public" field to submit a public offer.**

     A **Private offer** is submitted to one seller only. Only the seller will be informed of the offer and only the seller can accept the offer. **Click the "private"** field to submit a private offer. After that, you must specify which seller you want to submit the offer to by entering the seller's ID number. Remember, every seller maintains the same ID number throughout all periods of the experiment.

  b) Specify what **price** you want to offer. Enter your price in the "Your price" field. The price can range from 0 to 100 (whole numbers only).

  c) Specify what **quality** you desire. Enter this in the "Desired quality field". Quality can range from 0 to 10, where higher numbers are better (whole numbers only).

  After specifying the type of offer, the price, and the quality, click "OK" to submit it.

- On the left side of your screen, you will see the header "public offers," which displays all public offers made by buyers, including your own offer.
- In the middle of the screen, you can see all private offers that you have submitted in the current trading phase.
- **Each buyer can submit as many private and public offers as he wishes in each period.** Each offer that you submit can be accepted at any time during the trading phase.
- **In any given period, each buyer can conclude at most one trade.** Once one of your offers has been accepted, you will be notified which seller accepted which of your offers. This information will be displayed on the bottom right corner of your screen. At this point, all your other offers will be removed from the market and cancelled
- **In any given period, each seller can conclude at most one trade.** You will be continuously informed about which sellers have not yet accepted an offer. On the bottom right of the screen, you will see 7 fields, each field for one of the 7 sellers. Once a seller has accepted an offer, an "x" will

appear in the field next to the seller's identification number. You cannot submit private offers to a seller who has already concluded a trade.

- Once all 5 buyers have concluded a trade or after time has elapsed, the trading phase is over.
- No buyer is obliged to submit offers, and no seller is obliged to accept an offer.

2. Quality Determination Phase

- Following the trading phase, all **sellers** who have concluded a trade will determine the level of quality that they will supply to their buyers. **The product quality you asked for in your offer is not binding for your seller –** i.e. your seller can choose any quality he/she wants to from 1 to 10.

- While your seller is determining quality, you are asked to specify which quality you *expect* him/her to supply. In addition, we ask you to state how certain you are that the seller will actually deliver the quality you expect.

How Are Points (Income) Calculated?

Your Points

- If you do not conclude a trade during the trading phase, you will receive 0 points for that period.
- If one of your offers is accepted, your points depend on the price you offered and on the product quality. Your points for that period are determined as follows:

Your Points = 10*Product Quality – Price

- As you can see, the higher the product quality, the more points you earn. At the same time, the lower the price you paid, the more points you earn.
- Higher quality at lower prices means more points for you.

How do Sellers Earn Points?

- If a seller has not concluded a trade during the trading phase, he/she gains 5 points for that period.
- If the seller has accepted an offer, his/her income equals the price he/she receives minus the production costs he/she incurs. The income of a seller is determined as follows:

Points = Price – Production Costs

- As you can see, the higher the price, the more points a seller earns. At the same time, the higher the quality, the higher the production costs, which reduces points.

- How are production costs calculated? The higher the quality the seller supplies, the higher the costs. All sellers have the following cost table:

Quality	1	2	3	4	5	6	7	8	9	10
Cost	0	1	2	4	6	8	10	12	15	18

Points for all buyers and sellers are determined in the same way. **Each buyer can therefore calculate the income of his/her seller and each seller can calculate the income of his/her buyer.** Further, each buyer and seller is informed about the ID number of his/her trading partner in each period.

Please note that buyers and sellers can incur losses in each period (lose rather than gain points). These losses are subtracted from your points balance.

You will be informed about your points and the points of your seller in each period on an "**income screen**." The following information is displayed on this screen:

- the seller you traded with (ID number)
- the price you offered
- your desired quality
- the product quality you actually received from your seller.
- the points earned (lost) by your seller in this period.
- the points that you earned (lost) in this period.

Please enter all the information on the screen in the documentation sheet supplied to you. This will help you keep track of your performance across periods. After the income screen has been displayed, the period is over. Another period begins, starting with a trading phase. Once you have finished studying the income screen, please click "continue". All sellers also see an income screen displaying the same information.

Before we begin the experiment, we ask all participants to complete a questionnaire which will test your familiarity with the procedures. The experiment will not begin until all participants are completely familiar with all procedures.

In addition, we will conduct **2 trial periods of the trading phase** so that you can get accustomed to the computer. During the trial periods, no money can be earned.

Control Questionnaire

Please solve the following exercises completely. If you have questions, ask one of the experimenters. After all participants have answered the questions correctly, the experiment begins.

1. Suppose that you are a buyer and you did not make an offer during the trading phase. How many points do you earn for this period?

2. Suppose that you are a buyer and you offered a price of 30 and indicated a desired quality of 9. A seller accepts your offer and actually chooses a quality of 8. How many points did you earn for this period?

3. Suppose that you are a buyer and you offered a price of 10 and indicated a desired quality of 2. A seller accepts your offer and actually chooses a quality of 5. How many points did you earn for this period?

4. If a seller did not accept an offer during the trading phase, how many points does this seller earn for that period?

5. Suppose that you are buyer no. 3 in round 2. Does this mean that you will still be buyer no. 3 in round 3?

6. Suppose that you are a seller and you accepted an offer with a price of 30 and a desired quality of 9. You supplied an actual quality of 8. How many points did you earn for this period?

7. Suppose that you are a seller and you accepted an offer with a price of 40 and a desired quality of 2. You supplied an actual quality of 5. How many points did you earn for this period?

8. Suppose that you are seller no. 5 in round 5. Does this mean you were seller no. 5 in rounds 1-4 and will continue to be seller no. 5 in rounds 6-10?

If you have finished the exercises, we recommend that you look at the exercises and the solutions provided again. After this, you should think about the decisions you'll want to make to maximize your points.