Foundations of Relational Contracting: Experimental Evidence

Nisvan Erkal, Steven Y. Wu, and Brian E. Roe*

June 16, 2016

Abstract

We find considerable experimental support for several canonical predictions from relational contract theory, including a preference for informal agreements when third-party verification of performance is coarse, greater opportunistic behavior when discount factors decrease, and a tendency toward strategic ambiguity (Bernheim and Whinston, 1998). However, subjects inconsistently apply efficient punishments following a deviation, although behavior is consistent with a "less grim" version of Breitmoser (2015)'s semi-grim strategies. Since semi-grim strategies don't rule out conflict even after mutual cooperation, subjects appear to hedge against semi-grim strategic uncertainty by contracting for sub-optimal performance levels even when self-enforcement of optimal levels is possible.

KEYWORDS: relational contract, incomplete contract, endogenous incompleteness, informal incentives, experimental economics

JEL Classification: C73 C91 D86 J41 L14 L24 M52

---

*Erkal: University of Melbourne, Level 4, Faculty of Business and Economics Building 105, 111 Barry Street Carlton, Victoria 3010 Australia, n.erkal@unimelb.edu.au. Wu: Purdue University, Krannert Building, 403 West State Street West Lafayette, IN 47907, sywu@purdue.edu. Roe: The Ohio State University, 225 Agricultural Administration Building 2120 Fyffe Road Columbus, Ohio 43210, roe.30@osu.edu. We thank Tom Wilkening, Tim Cason, Vai-Lam Mui, Klaus Abbink and seminar participants at Purdue University, Monash University, and University of Sydney for helpful comments. Sharon Raszap and Amy Beth Corman provided excellent research assistance. Financial support from USDA-NIFA grant award no. 2010-65400-20430 is gratefully acknowledged. Roe recognizes support from the McCormick Program at The Ohio State University. This research is not the result of a for-pay consulting relationship. The experimental work was conducted under Purdue University IRB Protocol number 1103010650.

Relational contract theories capture important stylized features of observed contracting situations missed by complete contract theory. For example, contracts often omit easily verifiable performance factors (Scott, 2003) and contracting parties often trade repeatedly via informal self-enforcing agreements. While theories of relational contracts were first introduced several decades ago, interest in empirical studies has intensified recently. Empirical work using observational data is often constrained by difficulties in measuring key theoretical constructs such as discount factors, reservation payouts, one-shot deviation payouts from shirking, formal contracting alternatives, and other variables needed to specify self-enforcement and individual rationality constraints. Experiments can complement work based on observational data by allowing the researcher to directly specify a relational contracting model and perturb important variables such as discount factors, reservation utilities, costs, etc. in an environment free of unobserved institutional or social norm influences.

In this study, we use laboratory experiments to investigate relational and incomplete contract theory by employing an experimental design flexible enough to nest predictions consistent with several foundational theories (e.g., Telser (1980), Klein and Leffler (1981), Baker, Gibbons and Murphy (1994), MacLeod and Malcomson (1989), Bernheim and Whinston (1998), Schmitz and Schnitzer (1995) and Levin (2003) among others).

Our approach differs from most published experimental papers on relational contracts, which focus heavily on the role of social preferences and other intrinsic motivations. While such behavioral studies are important, we believe that experimental work based on canonical models is also needed to identify the strengths and weaknesses of standard models in order to more efficiently integrate behavioral and standard theories. The predictions that we focus on are based on a model with a rational self-interested principal who designs an optimal contract subject to individual rationality and self-enforcement constraints. To our knowledge, our study is the first experimental work to comprehensively test canonical predictions of relational contract theory.

Our experimental design starts from a first-best benchmark treatment where contracts can be made perfectly third-party verifiable. Additional treatments then deviate from first best. First, we alter the quality of formal contracts that can be written by introducing partial third-party verifiability, which mimics situations where a third-party can only verify crude performance outcomes such as whether a product is defective. Second, we create additional partial verifiability treatments where the discount factor varies, which affects the ability of contracting parties to self-enforce relational contracts. Infinitely repeated trading is implemented using a random continuation rule.

Another important feature of our design is that subjects assigned to be principals can choose from a large contract choice set including complete contracts and several types of incomplete contracts (e.g., fixed price contracts, discretionary bonus contracts, and pure bonus contracts). This allows for the examination of a wide range of predictions, including those emerging from the theory of strategic ambiguity (Bernheim and Whinston, 1998). Most previous experimental studies have imposed specific contractual forms (e.g., efficiency wages); thus, results from the received literature may not generalize to environments where subjects can choose optimal contractual forms.

Our main findings are as follows. First, we confirm a number of canonical predictions. These include: (a) when total pay does not meet individual rationality conditions or the promised discretionary bonus does not satisfy the agent's incentive compatibility condition, there is an increase in contract rejection or shirking; (b) with only partial contract enforcement, subjects shift towards relational contracts; (c) with a drop in discount factor, subjects shift to formal contracts in treatments featuring only partial contract enforcement; and (d) in the presence of imperfect verifiability, subjects largely choose discretionary bonus contracts rather than efficiency wage contracts, which is consistent with the theoretical optimality of discretionary bonus contracts in our model and the theory of strategic ambiguity of Bernheim and Whinston (1998).

Second, our results do not support the strict prediction that principals only use efficient punishments following contractual deviations. To put this into context, Levin (2003) suggests that, rather than terminating the relationship which destroys surplus and is therefore inefficient, there exists "strongly optimal" contracts that allow parties to continue with a relational contract but with terms restructured to hold the deviating party at his reservation payout.[1] While our results show that there is frequent use of inefficient punishments, we do find evidence to support more moderate versions of the theory. In particular, we find that the probability that subjects continue with a strongly optimal relational contract after a breach is consistent with the "semi-grim" strategies of Breitmoser (2015). In Breitmoser (2015), subjects play mixed strategies with a high probability of cooperation after mutual cooperation, intermediate probabilities of cooperation after one person defects, and low probability of cooperation after mutual defection. Our results generally follow this pattern, although punishments were "less grim" following deviations. This may be because, unlike Breitmoser (2015)'s subjects who played a repeated prisoner's dilemma game, our subjects' played a contracting game, so transfers and contract restructuring could be used in place of noncooperation to punish the deviator. We provide, to our knowledge, the first empirical investigation of parties' post-shirking strategies within the context of relational contracts.

1 Related Literature

Telser (1980) and Klein and Leffler (1981) are the first papers formalizing relational contracts. Both papers assume that third-party enforcement is not possible and show that the value of future exchanges can act as a private contract enforcement mechanism. During the next phase of theoretical advancements, MacLeod and Malcomson (1989) (MM), Baker, Gibbons and Murphy (1994) (BGM), Schmitz and Schnitzer (1995) (SS), Bernheim and Whinston (1998) (BW), and Levin (2003) delivered important insights into the structure of optimal relational contracts, the interaction between formal and informal contracts, and endogenous contractual incompleteness.[2]

Specifically, MM characterize the wage and performance outcomes that can be implemented by self-enforcing employment contracts in a model with sym-

[1]In addition, termination or reversion to formal contracts may not be renegotiation proof.

[2]See also Aghion and Holden (2011). In their survey article, they point out that "second generation" models of incomplete contracts tend to focus on relational contracting.

metric information. MM show that the optimal contract can take a variety of forms, ranging from high fixed price contracts (with threat of termination for poor performance) to discretionary bonus contracts. Levin (2003) characterizes optimal relational contracts under hidden information, moral hazard, and subjective performance evaluation. A key finding is that the optimal incentive contract with moral hazard resembles a one-step discretionary bonus contract. BGM and SS explore the interaction between formal and informal contracts. They find that formal and informal contracts act as substitutes if the default option is a formal contract rather than termination. Our experimental design is consistent with this environment and our results show that indeed relational contracts and formal contracts are substitutes.

The theory of strategic ambiguity described by BW suggests that, with verifiability imperfections, greater contractual incompleteness may enhance surplus by providing more discretionary latitude to use informal incentives. Our study can be seen as an experimental test of strategic ambiguity which is arguably one of the most important theories of incomplete contracting.

Interest in empirical testing of relational contracting theory has intensified. MacLeod (2007) discusses how relational contract theory can explain observed trading mechanisms, and Gil and Zanarone (2015) derive testable implications of relational contracting models and review recent empirical work, including recent contributions of Macchiavello and Morjaria (2015) and Antras and Foley (2015). We contribute to this nascent literature by providing experimental evidence for some of the canonical theoretical predictions.

The experimental literature on relational contracts has focused on the impact of behavioral theories on contracting outcomes rather than direct tests of canonical contract theory.[3] A work-horse model in this literature is the gift-exchange game, which is closely related to efficiency wages in that fixed prices but no bonuses are offered. Thus, in finitely repeated games, high unenforce-

[3]Note that when we refer to experimental work on relational contracts, we are focusing our discussion only on those experiments that involve repeat trading. We do not consider experiments that focus on one-shot transactions (e.g., Fehr, Klein and Schmidt, 2007) to be tests of relational contracting as they typically appeal to a different body of theory rather than classic self-enforcement constraints for addressing contracting inefficiencies.

able effort must be induced by reciprocity/fairness considerations (Casoria and Riedl, 2013; Gächter and Falk, 2002; Brown, Falk and Fehr, 2004, 2012). The main differences between these papers and ours is that our paper does not rely on social preference arguments but is based instead on a classic infinitely repeated game model where we explicitly vary the strength of self-enforcement constraints to generate different equilibrium predictions.

The experimental paper by Sloof and Sonnemans (2011) tests both behavioral theory and a canonical prediction concerning the interaction between explicit and relational incentives. Thus, there is a small overlap between their paper and ours. Both papers find that weaker explicit contracts can support stronger relational contracts. This consistency in results holds despite substantial design differences between the two papers, which suggests that the finding that weaker explicit incentives support informal incentives is robust.[4] A key difference is that we focus on a wider array of canonical predictions beyond the tradeoff between explicit and informal contracts.

2 Theoretical Predictions and Implications

2.1 Model setup

We describe a simple model that can conceptualize many of the standard predictions from the relational contracting literature and form the basis for our experimental design. Our purpose is not to derive new theory but rather to provide a parsimonious unifying framework for many canonical results that span several papers in the literature. Such a simple, unifying framework serves the dual purpose of providing clear intuitive predictions and facilitating laboratory implementation where simplicity is not only a virtue, but a necessity.

Due to space constraints, we will only provide an abbreviated description

[4]For example, Sloof and Sonnemans (2011) use a reduced form contracting environment where they impose a restricted set of contracts (trust games) on subjects. Our design is based on a fully specified principal-agent model that imposes very few restrictions on contractual form; i.e., contract structure emerges endogenously subject to incentive compatibility constraints, individual rationality constraints, and enforcement limits. In addition, in their paper, it is the agent who designs the contract whereas in our model, it is the principal.

of our model, focusing on the empirical implications that follow from canonical predictions. We refer interested readers to the Appendix where the model is fully described along with propositions, proofs, and detailed discussions about how the empirical implications connect to the propositions.

Assume a principal contracts with an agent to produce a unit of a good for which quality matters. For simplicity, we abstract from asymmetric information, so our environment is similar to MM where the key friction is the absence of third-party enforcement. The agent's obligation is to deliver quality $q \geq Q$ where Q refers to the quality level specified in the contract and q refers to the actual quality delivered. The principal's obligation is to pay $w \geq W$ where w is actual payment and W is the payment specified in the contract. w can consist of a base price p and bonus payment b, so we write $w = p{+}b$. Similarly, we write $W = P + B$ for the contractually specified payments. Since P is a fixed, non-contingent payment, $p{=}P$ by default.

The principal's and agent's payoffs are $\pi_P = r(q){-}p{-}b$ and $\pi_A = p{+}b{-}c(q)$ where $r(q)$ and $c(q)$ are differentiable such that $r'(q) > 0$, $r''(q) \leq 0$, $c'(q) > 0$ and $c''(q) \geq 0$, $\forall q \in [\underline{q}, \overline{q}] \subset \mathbb{R}_+$. All else equal, the principal prefers higher quality and lower payments, and the agent prefers the opposite. The reservation payoffs for the principal and agent are $\overline{\pi}$ and $\overline{u}$, respectively.

2.2 Nesting Formal and Relational Contracts

We assume limited third-party verifiability where a third-party is able to detect whether the good achieves some coarse, discrete level of quality but cannot detect more refined gradations in quality. Limited third-party verifiability allows for imperfections in performance measurement in the spirit of BGM, but it conceptualizes the issue in a simpler one-dimensional framework that facilitates experimental implementation. Moreover, in practice, many products receive discrete quality certifications that are neither completely unenforceable by a third-party nor enforced to highly refined quality grades. Thus, our setup better matches stylized observations while allowing for a nesting of both formal and informal contracts in a parsimonious framework.

To model partial verifiability, we partition the quality space $[\underline{q}, \overline{q}] \in \mathbb{R}_+$ into $[[\underline{q}, q^d), [q^d, \overline{q}]]$ where q^d is a quality threshold that can be feasibly verified by a third-party. Thus, a third-party can verify whether $q \in [\underline{q}, q^d)$ or $q \in [q^d, \overline{q})$ This implies a contractible set, $\underline{C} = \{\underline{q}, q^d\}$.[5]

Enforcement imperfections do not preclude the possibility of writing formal/complete contracts, though imperfections do limit the set of available complete contracts.[6] The complete contract can either specify state-contingent prices $\underline{P}$ and P^d for each contractible quality realization, or the principal can specify $Q = q^d$ in exchange for a fixed P. We will refer to the latter as a **simple contract**. In the former case, a third-party enforces the contingent payments $\underline{P}$ and P^d whereas in the simple contract, $Q = q^d$ and P are directly enforced. In either case, all variables are third-party enforceable since they are either in the contractible set or depend only on variables in the contractible set. If the contingent payments $\underline{P}$ and P^d are chosen in an incentive compatible manner to implement $Q = q^d$, then the two types of contracts are outcome equivalent. Thus, for simplicity, we will focus only on simple contracts.

To model endogenous incompleteness, we denote π^f and u^f as the payoffs obtained from the "best" complete contract for the given enforcement technology; i.e., the formal contract that yields the highest joint surplus under the enforcement technology. In our case, if the first best quality level is such that $q^* > q^d$, then a formal contract specifying q^d would dominate one specifying $\underline{q}$. Since there are only two contractable quality levels, the contract specifying q^d is the best complete contract. Denote Q^f as the best contracted quality

[5]No other quality level is verifiable; hence, the agent will choose $q = q^d$ even if a contract calls for $Q > q^d$ and will choose $q = \underline{q}$ if the contract calls for $\underline{q} < Q < q^d$.

[6]A formal contract must be a complete contract in that a complete state-contingent plan governs performance. Therefore, all obligations of both parties are fully specified for all contingencies in the initial contract. Moreover, the contract is third-party enforceable so that neither party can shirk. This implies that no party has ex post discretionary latitude to deviate from the initial contract. One can view the presence of ex post discretion to deviate as being synonymous with an incomplete contract. This implies that the contract would have to be self-enforcing through an informal agreement.

level.[7] Denote surplus as $S(q) = r(q) - c(q) - \overline{u} - \overline{\pi}$. We define

$$k = S(q^*) - S(Q^f) \tag{1}$$

to be the loss in efficiency from using a formal contract in the presence of verifiability imperfections. Note that when a third-party can verify every quality level, then $k = 0$ since $Q^f = q^*$. Like BGM, our model can nest formal and informal contracts. Unlike BGM, we have a single performance measure rather than separately defining objective and subjective measures. This setup eases experimental implementation since subjects track fewer variables.

2.3 Optimal Contracting

Consider a principal-agent model of repeat trading under the imperfect enforcement technology specified above. Define a binary variable $\alpha \in \{0, 1\}$ where α equals 1 if $u^f + \pi^f \geq \overline{u} + \overline{\pi}$ and 0 otherwise. That is, $\alpha = 1$ if joint profits from the best complete contract exceeds joint reservation payoffs. The stage-game timeline follows the typical principal-agent sequence:

1. Principal offers a contract–a price/bonus/quality triplicate, (P, B, Q).

2. The agent accepts or rejects. If rejected, the parties default to the best formal contract if $\alpha = 1$ and to reservation payoffs if $\alpha = 0$.

3. If accepted, the agent chooses actual quality q.

4. The principal observes q and chooses actual bonus b. The promised fixed payment, P, is also made.[8]

A relational contract is an infinite repetition of the above stage-game so that in each period t and for each history up to t, the relational contract describes the sequence (1)-(4). Moreover, the relational contract is self-enforcing if it describes a subgame perfect equilibrium of the infinitely repeated game.

[7] In our example $Q^f = q^d$.

[8] P is always third party enforceable because it is not contingent on quality.

In addition, Levin (2003) and Halac (2012) show that, with symmetric information, there exist stationary contracts that are optimal in that the same (optimal) contract is offered in every t.[9] Letting δ be the discount factor and multiplying the payoffs by $1-\delta$ to express them as per-period averages, the principal's contract design problem is:

$$\max_{Q,P,B}(1-\delta)\left[r(Q)-P-B\right]+\delta V\left[C\right] \quad s.t. \tag{2}$$

$$(1-\delta)\left[r(Q)-P-B\right]+\delta V\left[C\right]\geq \alpha\pi^{f}+(1-\alpha)\overline{\pi} \tag{3}$$

$$(1-\delta)\left[P+B-c(Q)\right]+\delta U\left[C\right]\geq \alpha u^{f}+(1-\alpha)\overline{u} \tag{4}$$

$$(1-\delta)\left[r(Q)-P-B\right]+\delta V\left[C\right]\geq(1-\delta)\left[r(Q)-P\right]+\delta\left[\alpha\pi^{f}+(1-\alpha)\overline{\pi}\right] \tag{5}$$

$$(1-\delta)\left[P+B-c(Q)\right]+\delta U\left[C\right]\geq(1-\delta)\left[P-c(\underline{q})\right]+\delta\left[\alpha u^{f}+(1-\alpha)\overline{u}\right] \tag{6}$$

Constraints 3 and 4 are the individual rationality (IR) constraints and 5 and 6 are the self-enforcement (SE) constraints. $V(C)$ and $U(C)$ can be understood as follows: let Γ denote the set of feasible contracts, which can be partitioned as $C\cup F=\Gamma$ and $C\cap F=\emptyset$. Then, either $(P,B,Q)\in C$ or F, where "C" denotes relational contracts that satisfy contraints 3-6, and "F" denotes "formal" (i.e., complete) contracts that only satisfy the IR constraints. Thus, $V(C)$ and $U(C)$ are the flow payoffs for the principal and agent, respectively, from the optimal self-enforcing relational contract $(P,B,Q)\in C$. Due to stationarity, the same contract is offered every t, so the principal's contract design problem becomes essentially a static optimization problem.

[9]Nonstationary contracts arise primarily in the context of private information where one has to model relational dynamics due to the revelation of private information over time (e.g., see Halac, 2012 or Yang, 2013). It is important to point out that nearly all experiments involve some dynamics simply because subjects learn how to play the game. Hence, researchers typically treat predictions from stationary symmetric information games as theoretical benchmarks that subjects should converge to after sufficient learning. The actual dynamics that lead to convergence is typically not of theoretical interest and early period departures from theoretical benchmarks are treated as noise that can be reduced with subject experience.

Solving the above model yields an optimal stationary relational contract. In addition, a number of propositions and corollaries follow which we state in detail in the Appendix. These propositions and corollaries lead to a number of empirical implications which we discuss in the following section.

3 Empirical Implications and Experimental Design

3.1 Empirical Implications

The first empirical implication follows from the fact that the optimal contract implements some $\tilde{Q}$ that is less than or equal to first best quality, Q^* using a discretionary bonus that simultaneously satisfies both the agent's and principal's SE constraints, combined with a base price, P, that ensures that both parties' IR constraints are met. The principal's and agent's SE constraints 5 and 6 can be combined and rewritten as:

$$\delta\left[r(Q) - P - \alpha\pi^f - (1-\alpha)\overline{\pi}\right] \geq B \geq (1-\delta)\left[c(Q) - c(\underline{q})\right] - \delta\left[P - c(Q) - \alpha u^f - (1-\alpha)\overline{u}\right] \tag{7}$$

Empirical Implication 1. *Discretionary bonuses, $B(\tilde{Q})$, that violate the l.h.s. of 7 are non-credible and will lead to increased contract rejection.*[10] *$B(\tilde{Q})$ that violate the r.h.s. of 7 will lead to increased shirking by agents. Promised total payments that do not satisfy the agent's IR constraint will increase contract rejection rates.*

Levin (2003)'s Corollary 1 (p. 841) points out that, because optimal stationary contracts can be constructed to split the surplus in any way the parties desire (subject to IR constraints), the parties can continue with a relational

[10]In principle, $B(\tilde{Q})$ that violate the l.h.s. of 7 should also increase shirking on the bonus by the principal. However, since the principal both sets $B(\tilde{Q})$ and makes the decision on actual bonus b, this is plagued by endogeneity problems. A principal who specifies a non-credible $B(\tilde{Q})$ may have no intention of honoring the bonus in the first place so promised bonus and actual bonus are jointly determined.

contract even following a deviation. Levin (2003) shows that, following any history, including those that are off-the-equilibrium path (i.e., a deviation), there is a family of **strongly optimal** relational contracts that implement $\tilde{Q}$ while delivering different payoff distributions. Thus, one can always construct an off-the-equilibrium path contract that continues to implement $\tilde{Q}$, while holding the deviator to the payoff he would have received had the parties reverted to a formal contract or termination. In other words, the deviator can be punished as severely as termination of the relational contract, but without destroying surplus and without also punishing the non-deviator. Such a contract does not destroy surplus since surplus is higher under $\tilde{Q}$ than under Q^f or termination and is therefore renegotiation proof. In short, continuing with a relational contract is optimal regardless of whether the parties have deviated or not in the previous period.

Empirical Implication 2. *Following a deviation, the parties should respond with the most efficient punishment mechanism, which is to continue with a relational contract, but with terms adjusted to hold the deviating party to his formal contract payoff, or reservation payoff, whichever is higher.*

Next, we look at the impact of verifiability on relational contracting. For a more intuitive look at self-enforcement, solve 7 for δ which yields:

$$\delta \geq \underline{\delta}(Q) = \frac{c(Q) - c(\underline{q})}{r(Q) - c(\underline{q}) - \alpha\left[\pi^f + u^f\right] - (1-\alpha)\left[\overline{\pi} + \overline{u}\right]} \tag{8}$$

$$= \frac{c(Q) - c(\underline{q})}{r(Q) - c(\underline{q}) - \alpha\left[r(Q^f) - c(Q^f)\right] - (1-\alpha)\left[\overline{\pi} + \overline{u}\right]} \tag{9}$$

$\underline{\delta}(Q)$ is the threshold for the incomplete contract to be self-enforcing, and it depends on Q, where a higher Q raises the threshold making self-enforcement more difficult. Hence, this can limit the quality that can be implemented.

Empirical Implication 3. *A decrease in δ weakly decreases Q that the principal contracts for and/or increases the use of formal contracts.*

The threshold also depends on the payoffs u^f and π^f, which in turn, depends on the efficiency loss from imperfect verifiability. Thus, self-enforcement and third-party enforcement interact; i.e. suppose Q^f is the enforceable quality that yields the highest joint surplus among all contractible quality levels. A complete contract (Q^f, P^f) yields payoffs $\pi^f = P^f - c(Q^f)$ and $u^f = P^f - c(Q^f)$. These payoffs can be substituted in (8) to get (9). As k in (1) tends toward zero, third-party verifiability improves. This, in turn, increases the joint profit $r(Q^f) - c(Q^f)$ which weakly raises the threshold for self-enforcement 8.[11] In short, an improvement in enforcement technology should cause some relational contracts to be replaced by complete contracts.

Empirical Implication 4. *Moving from partial verifiability to full verifiability leads to more formal/complete contracts.*

The above implication 4 is related to the theory of *strategic ambiguity* of BW and to the substitutability between formal and informal contracts of BGM. BW show that, in the presence of verifiability imperfections, parties may deliberately eschew formal contracts so that they can use discretionary flexibility to punish and reward non-verifiable performance.

Another BW insight is that, given that contracts must be incomplete, it may be optimal for parties to *increase* the degree of incompleteness. Intuitively, under an incomplete contract, the agent has ex post discretionary latitude to shirk. Thus, the principal may also leave herself with discretion via a discretionary bonus contract so that she can adjust pay in response to the agent's action. Such a contract is less complete than a fixed-price contract because the fixed-price contract locks down the principal's obligations. While fixed price contracts are commonly invoked in the literature under the assumption that parties to a relational contract use efficiency wages or repeat purchase mechanisms (Klein and Leffler, 1981; Shapiro and Stiglitz, 1984; Brown, Falk and Fehr, 2004), they are not consistent with the theory of strategic ambiguity.

Empirical Implication 5. *We should not observe efficiency-wage type fixed-price contracts as the preferred contractual form for relational contracts.*

[11] We say weakly because if $\alpha = 0$, then the threshold does not change until complete contracts joint surplus exceeds joint surplus from the reservation payoffs, triggering $\alpha = 1$.

3.2 Experimental Design

Our experimental design is based on the above contracting model. We impose specific parameters and functional forms, which are chosen to obey the curvature assumptions of the above model to minimize loss of generality. In the experiments, we refer to the principals as "buyers" and agents as "sellers."

A crucial design feature is that buyers can endogenously choose contractual form subject to exogenously imposed verifiability limits. To achieve this, we specify sellers' action space as $q \in \{1, 2, ..., 15\}$. We define two enforcement technologies which represent a major treatment variation:

1. **Technology E**: Perfect enforcement technology allows a third-party to verify and enforce every single quality level in $\{1, 2, ..., 15\}$ so that the contractible set is $\underline{C}_E = \{1, 2, ..., 15\}$.

2. **Technology PE**: Partial enforcement technology partitions the quality space as $\{\{\underline{q}, ..., q^d - 1\}, \{q^d, ..., \overline{q}]\}\} = \{\{1, 2, 3, 4\}, \{5, 6, ..., 15\}\}$ where $q^d = 5$. The contractible set is thus $\underline{C}_{PE} = \{1, 5\}$.

Technology E provides perfect quality grading whereas Technology PE allows a third-party only to arbitrate on whether the product was defective (i.e., below 5). Thus, even under PE, the parties can write a complete contract that conditions on whether the product is defective.

Buyers can endogenously structure complete contracts under E and PE as follows. Within each stage-game, each buyer can (but is not required to) offer a contract, (P, B, Q) where $P \in \{0, 1, ..., 200\}$ is a fixed price, $B \in \{0, 1, ..., 200\}$ is a discretionary bonus, and $Q \in \{1, 2, ..., 15\}$ is the buyer's requested quality level. Buyers can endogenously specify "simple" complete contracts by specifying $Q \in \underline{C}_E$ or $Q \in \underline{C}_{PE}$, depending on the treatment, along with a fixed price P and then clicking a "binding" option on the computer screen. When binding is checked, neither party has ex post discretionary latitude to deviate as the computer enforces P and Q. A discretionary bonus B is redundant since it plays no incentive role as the seller cannot deviate from Q. Apart from these restrictions, we impose no other structure on contracts; i.e., subjects can endogenously specify complete contracts, as well as a range of incomplete contracts seen in the literature, including gift-exchange/efficiency wage ($P > 0, B = 0$), discretionary bonus ($P > 0, B > 0$), and pure bonus ($P = 0, B > 0$). Specifying an incomplete contract only requires the buyer to check the "discretionary" box rather than the "binding" box. When discretionary is checked, Q and B are not enforced by the computer.

The following summarizes the sequence of events in a stage-game.

1. **Proposal phase**-buyer can offer a single contract (P, B, Q) to seller. The seller can accept or reject; hence an IR constraint is active.

2. **Quality phase**-seller chooses q if Q is not binding.

3. **Payment phase**-buyer chooses actual b (if B was in the contract).

Under binding contracts, there are no **Quality** or **Payment** phases since neither party can deviate from the initial contract. Stage-game payoffs are $\pi = 12q - P - b$ and $u = P + b - (q^2)/2$ for the buyer and seller, respectively. Sellers are provided with Table 1 so that they can quickly calculate costs. Reservation payoffs are $\overline{\pi} = \overline{u} = 15$, which are triggered if either the buyer does not offer a contract or the seller rejects a contract. First best is realized at $q = 12$ which yields joint payoffs of 72 and exceeds the joint payoffs of 30 if parties do not contract. If $q < 3$, then joint profit is below the joint outside option payoffs of 30, making it risky for the parties to engage in contracting. Additionally, in the PE treatments, the best contractible quality is $Q^f_{PE} = 5$ which yields joint payoffs of 47.5 and exceeds joint outside option payoffs. In the E treatment, the first best level is in the contractible set, so $Q^f_E = 12 = Q^*$.

We follow the typical approach of implementing an infinitely repeated game

Table 1: Seller's Cost

Quality	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Cost	1	2	5	8	13	18	25	32	41	50	61	72	85	98	113

using a random continuation rule (e.g., Bó, 2005). Specifically, we exogenously form buyer-seller pairs where the pair can trade with each other for a random number of stage-games. In each period, there is δ probability that the same buyer and seller will trade with each other again the next period. This allows a second treatment variation:

1. **0.8 treatment**: $\delta = 0.8$
2. **0.5 treatment**: $\delta = 0.5$

Self-enforcement is obviously stronger in the 0.8 treatment because there is a larger probability that the parties will trade next period. We refer to the repeated game for each buyer-seller pair as a *supergame*. Thus, the supergame is expected to last five(two) periods for the 0.8(0.5) treatment.[12]

Table 2: Treatments

	$\delta = 0.50$	$\delta = 0.80$
Perfect enforcement (E)		two sessions
Partial enforcement (PE)	three sessions	three sessions

Treatment variations are summarized in Table 2. We run the E treatment under $\delta = 0.80$ only because, by Empirical Implication 4, when enforcement is perfect, complete contracts should be used regardless of δ. Since incomplete contracts are more likely to be seen with $\delta = 0.80$, if subjects use complete contracts under E0.80, then they will use complete contracts under E0.50.

All interactions between subjects occur via computers and subjects identify each other by assigned ID numbers that are not associated with actual identities. Once subjects are seated, the program randomly assigns half the subjects to be "buyers" and the other half to be "sellers." Roles were fixed for the duration of the experiment. Subjects are then read instructions and answer some control questionnaires to ensure understanding. We subsequently conduct two trial periods to acclimate subjects to the trading platform. ID numbers are suppressed during the trial periods. Once the live rounds begin, each buyer is exogenously matched to a seller to play a supergame which consists of a sequence of repeated stage-games until randomly terminated. No subjects are matched for more than one supergame (stranger matching).

[12]The expected number of periods is $\frac{1}{1-\delta}$.

The experiment ends when one of two conditions occurred: (1) All possible supergame matches are exhausted and the last pairing randomly terminates; (2) If all pairings are not exhausted and the subjects play at least 18 periods (across all supergames) in the $\delta = 0.8$ treatment or at least 20 periods in the $\delta = 0.50$ treatment, then they are in their last supergame and the experiment ends when that supergame terminates. These long sessions ensure there is adequate opportunity for learning. We recruited either 20 or 22 subjects per session for the $\delta = 0.5$ experiments and either 16 or 18 subjects per session for the $\delta = 0.8$ experiments.[13]

Experiments were conducted in the Vernon Smith Experimental Economics Laboratory (VSEEL) at Purdue University, a lab with an explicit no deception policy. The subject pool consisted of undergraduate students in the VSEEL subject database. Subjects may have participated in other experiments but not our specific treatments. Nine sessions involving 170 subjects were conducted under an approved IRB protocol. All payoffs are given in points, which accumulate across periods, and converted into U.S. dollars at the rate of 30 points=$1. This method of payment is common in repeated game experiments (e.g., Bó, 2005). Average pay exceeded 25 USD per-session, with a range from $15 to $38, which includes a $5 show-up fee. The average session lasted about three hours, including instructions, questionnaire, trial periods, post experimental payouts and post experimental demographic questionnaire. Average hourly payouts match hourly rates of other experiments conducted in the same lab. All experiments were programmed with Z-tree (Fischbacher, 2007).

[13]We recruited more subjects for the $\delta = 0.5$ treatments because the expected length of supergames are shorter. Thus, we would likely exhaust matches more frequently in the $\delta = 0.5$ sessions if we did not recruit more subjects. The differences in group size should not create an imbalance in group reputation effects since we implemented stranger matching.

4 Results

4.1 Empirical Implication 1: Credibility of the Discretionary Bonus

Empirical Implication 1 states that when B is so large that it violates the l.h.s. of inequality 7, then B is non-credible and the agent will reject the contract. Conversely, when B is so small that it breaches the r.h.s. of 7, then it lacks the power to induce the agent to deliver $q \geq Q$. Finally, if promised profit under the contract does not satisfy the agent's IR constraint, the agent will reject the contract.

We use the upper and lower bounds in inequalities 7 to create two dummy variables: *noncredible B* takes a value of "1" if a contract contains B greater than the upper bound and *nonIC B* equals "1" if B is less than the lower bound. We also created a dummy *IR-satisfied* that equals "1" if the promised profit to the seller under a contract exceeds the seller's reservation payoff.

The first two regressions in Table 3 are linear probability models (LPM) of the seller's rejection decision (=1 if reject, 0 otherwise). Regression (2) adds seller fixed effects, since unobserved seller heterogeneity could create selection effects into certain types of contracts so that the error term may be correlated with the contract dummies. Regression (2) also includes a 1-memory cooperation dummy that equals "1" if the parties engaged in and honored (i.e., $b \geq B$ and $q \geq Q$) a relational contract in the previous period. This dummy is included to account for the possibility that a seller might form and update beliefs about a buyer's actions.

The probability of rejection declines when the IR constraint is satisfied (-0.365, $p < 0.01$ in regression (1) and -0.289, $p < 0.01$ in regression (2)), which is consistent with the theory. These results appear to be robust as the coefficient estimates and significance do not vary greatly across the two specifications. The coefficients for *noncredible B* are positive, but they are significantly different from zero only in regression (1) and only at the 10% level of significance. Thus, there is only tentative evidence that sellers are

Table 3: LPM Estimates (PE0.50 and PE0.80 data pooled)

	Binary Dependent Variable			
	(1)	(2)	(3)	(4)
	Seller Reject=1	Seller Reject=1	Seller Shirk=1	Seller Shirk=1
noncredible B(dummy)	0.06*	0.058		
	(0.036)	(0.058)		
nonIC B(dummy)			0.264***	0.471***
			(0.10)	(0.131)
IR-satisfied(dummy)	-0.365***	-0.289***		
	(0.083)	(0.051)		
1-memory cooperation dummy		-0.239**		-0.26***
		(0.067)		(0.061)
PE0.80(dummy)	-0.155***	0.0135	-0.20*	-0.999***
	(0.037)	(0.064)	(0.083)	(0.021)
Period	0.016	0.049*	-0.014	-0.0002
	(0.012)	(0.021)	(0.0095)	(0.027)
Period^2	-0.0005	-0.0009	0.0007	-0.0001
	(0.0005)	(0.0008)	(0.0005)	(0.001)
Constant	0.46***	0.137	0.774***	1.00***
	(0.065)	(0.105)	(0.083)	(0.049)
Seller fixed effects	No	Yes	No	Yes
Observations	560	291	382	189

-Robust standard errors clustered on sessions are reported in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

forward looking enough to reject non-credible bonus offers.

Regressions (3) and (4) examine the seller's shirk decision (dependent variable=1 if $q < Q$). The estimated coefficients for *nonIC B* are positive and significant (0.264, $p < 0.05$ in regression (3); 0.471, $p < 0.001$ in regression (4)) suggesting that incentive compatibility motivates sellers to honor their agreements. These results appear to be robustly consistent with theory.

To summarize, our data largely supports Empirical Implication 1, but the non-credibility of B has only a weak impact on the seller's decision to reject a contract. Our results suggest that the theoretical SE and IR constraints needed to solve for optimal contracts have important empirical relevance.

4.2 Empirical Implication 2: Efficient Punishment

Empirical Implication 2 states that, following a deviation by either party, the most efficient punishment mechanism is for the buyer to continue to offer a relational contract but adjust the terms so that rent is shifted away from the party that deviated. Switching to a formal contract or terminating the relationship are inefficient punishments.

Figure 1 suggests that when one or both parties shirked in the previous period, the buyer only offers a relational contract about 36% of the time. This appears to contradict Empirical Implication 2.

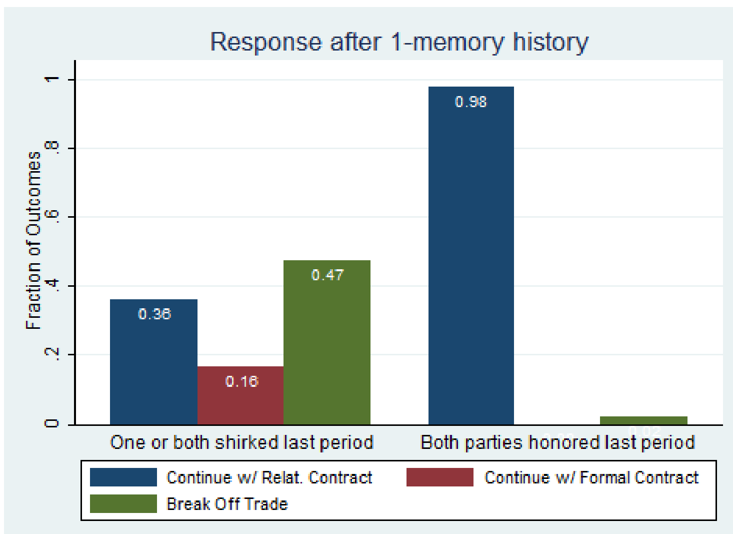

Figure 1: *Buyer response after 1-memory cooperation/non-cooperation (combined PE0.50 and PE0.80 data, all rounds)*

For more detail, we partition the 1-memory state space into four states: Both parties honor (H,H); buyer honors but seller shirks (H,S); buyer shirks but seller honors (S,H); and both shirk (S,S). We estimate LPMs of a buyer offering a relational contract as a function of four state dummy variables, one for each state, with multi-level session-buyer-seller random effects and robust-standard errors clustered on sessions (Table 4).[14]

[14] Breitmoser (2015) argues that controlling for unobserved heterogeneity is important since observations from cooperative states are more likely to come from cooperative types.

Table 4: Prob. of Relational Contracting After 1-Memory Histories

	(1) PE50 Treatment	(2) PE80 Treatment	(3) PE50+PE80
Both honored (H,H)	0.80^{***}	0.996^{***}	0.93^{***}
	(0.139)	(0.004)	(0.042)
Buyer honored (H,S)	0.333^{***}	0.400^{***}	0.362^{***}
	(0.096)	(0.116)	(0.082)
Seller honored (S,H)	0.286^{***}	0.632^{***}	0.444^{***}
	(0.064)	(0.145)	(0.086)
Neither honored (S,S)	0.348^{***}	0.523^{***}	0.428^{***}
	(0.036)	(0.077)	(0.054)
Observations	94	127	221

-Robust standard errors clustered on sessions are reported in parentheses.

-Linear probability models estimated with random effects at the session-buyer-seller levels.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

The strict version of Empirical Implication 2 suggests that the probability of a buyer continuing with a relational contract should be close to "1" for all 1-memory states. We can see clearly that this does not hold. While the estimated probability of the principal offering a relational contract ("cooperation") is highest after mutual cooperation (H,H) (0.80 for the PE0.50 data and 0.996 for the PE0.80 data), the probabilities drop off significantly after at least one party shirks. For the PE0.50 data, the estimated probabilities for the three shirking states range from 0.286 to 0.348. The estimated probabilities are higher in PE0.80 (0.40 to 0.632), but are still far below 1.

Thus, our results do not support the strong prediction that subjects always use the most efficient punishment mechanism. At the same time, the evidence seems to support more moderate versions of the theory for two reasons. First, the higher estimated probabilities of continuing with a relational contract in PE0.80 versus PE0.50 suggests that when the loss from using inefficient punishments is greater, subjects use inefficient punishments less often. Second, subjects' average strategies follow a pattern that resembles a "less-grim" version of the "semi-grim" strategies noted by Breitmoser (2015) (B15) in the context of a repeated prisoner's dilemma (PD) games.

B15 found that repeated game strategies are well described by 1-memory Markov "semi-grim" mixed strategies where parties cooperate with high probability after mutual cooperation, defect with high probability after mutual defection, and randomize with intermediate probability when only one player has defected. Unlike PD games, however, contracting offers the possibility of continued cooperation after a defection because parties are able to make transfers through pay adjustments to reward and punish rather than resort to non-cooperation. Whereas B15 finds about 10% cooperate after (S,S), we find that a relational contract will be offered with about a 35% chance in PE0.50 and a 52% chance in PE0.80. The higher probability of cooperation in our experiments may suggest that parties learn that switching away from a relational contract is a less efficient means of punishment.

Table 5 presents average contract terms for relational contracts used after 1-memory states. For the PE0.80 treatment, there is a clear pattern of contract term adjustments across 1-memory states. Using (H,H) as a benchmark, note that after (H,S), promised profit to shirking sellers drops dramatically.[15] If instead, buyers shirk but sellers honor (S,H), buyers do not seem to reward sellers with higher promised profits (35.57 vs 36.67), but do offer higher fixed payments, P, and lower discretionary bonuses, B. This suggests that buyers try to reduce strategic uncertainty faced by sellers. Intuitively, if a seller honors the contract ($q \geq Q$) while the buyer shirks on the bonus ($b < B$), then the buyer may have to provide assurances in the next contract or the seller will reject. Finally, when both parties shirk (S,S), buyers respond by offering contracts that promise less profit to sellers (25.71 vs 36.67), but also provide them with more security by raising P and lowering B. Intuitively, when both parties failed to cooperate in the prior period, sellers may need more guarantees to continue with such a strategically uncertain relationship.

In the PE0.50 treatment, there also appears to be adjustments after 1-memory states. Like the PE0.80 treatments, buyers appear to offer higher P

[15]Promised profits are what the parties would earn if both parties honored the contract.

Table 5: Relational Contract Terms After 1-Memory Histories

	(H,H)	(H,S)	(S,H)	(S,S)
PE80 treatment(Means)				
B	54.48	12	24.14	35.71
P	36	42.5	48.71	42.55
Q	10.6	10.17	8.43	9.97
Promised Seller Profit	36.67	1	35.57	25.71
Promised Buyer Profit	31.26	67.5	28.29	41.35
N	42	6	7	31
PE50 treatment(Means)				
B	59.5	15	58.25	51
P	20	49	38.75	45.42
Q	9.75	13	9	10.29
Promised Seller Profit	31.75	-21	52	38.79
Promised Buyer Profit	37.5	92	11	27.08
N	4	2	4	24

in (S,H) and (S,S) to provide sellers with more security. At the same time, buyers maintain high B in (S,H) and (H,H) perhaps to motivate sellers to deliver high q which is more difficult when $\delta = 0.5$. The net effect is that sellers are offered high promised profits in (S,H) and (H,H) in the PE0.50 treatment, perhaps because high upfront payments (P) and high discretionary bonuses (B) are needed to ensure their participation and delivery of high quality when self-enforcement is weaker. Nonetheless, these patterns should be interpreted with caution since there were so few observations outside the (S,S) state.

To summarize, our results do not support the strict version of the theory that predicts that subjects will use "strongly optimal" relational contracts after deviations. However, our results seem to support a more moderate version of the theory that is consistent with a "less grim" version of B15's semi-grim strategies. Compared to B15's prisoner's dilemma semi-grim results which suggests a mixed strategy profile of (0.9,0.3,0.3,0.1) for states that are analogous to our ((H,H),(H,S),(S,H),(S,S)) states, we find a "less-grim" mixed strategy profile of roughly (0.9, 0.4, 0.4, 0.4).[16] Thus, "cooperation" (i.e., the continued use of relational contracts) in our experiments is slightly higher after one party deviates and significantly higher after mutual shirking relative to B15's findings. The ability to re-adjust contract terms to either reallocate promised profits or to alter the degree of strategic uncertainty can explain the higher cooperation rates under contracting as opposed to prisoner's dilemma games.

4.3 Empirical Implication 3: Impact of δ

Empirical Implication 3 suggests that, when the discount factor drops, the principal will contract for a lower Q or switch to a formal contract.

Under our experimental parameters, the maximimum self-enforcing Q in the PE0.50 treatment is $Q = 8$. However, $Q = 8$ does not yield substantially more surplus than a formal contract that guarantees $Q = 5$. Thus, it is possible that some parties opt to use formal contracts in the PE0.50 treatment, though this is an empirical question. On the other hand, under $\delta = 0.80$, the parties should theoretically be able to contract for $Q = 12$, the first-best level. Hence, when δ drops from 0.80 to 0.50, we should observe:

1. More formal contracts used in the PE0.50 treatment relative to the PE0.80 treatment.
2. A reduction in Q for those who still use relational contracts in PE0.50.

Figure 2 plots the average fraction of complete contracts for each period by treatment, using data from all sessions. Since the number and lengths of supergames differed by session and treatment, Figure 2 cannot distinguish the fraction of complete contracts across supergames. Nonetheless, the figure displays how play evolved as subjects gained experience.

[16]For robustness, we also estimated the profile using fixed effects LPM, which yielded a profile of roughly (0.85, 0.4, 0.5, 0.45). These values are close to what we obtain under random effects and does not fundamentally alter our conclusions.

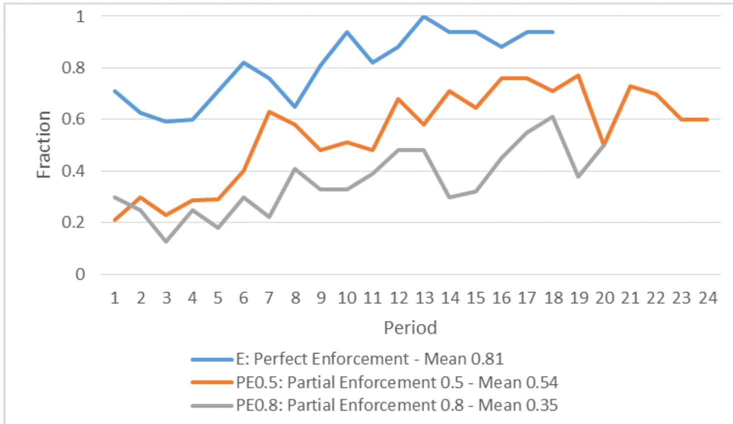

Figure 2: *Average fraction of complete contracts across periods for all sessions.*

Note that the fraction of complete contracts is higher in the PE0.50 treatment versus the PE0.80 treatment, which is consistent with the theory. The overall mean fraction of complete contracts is 0.54 in PE0.50 and only 0.35 in PE0.80 and the gap persists through most of the periods with the exception of an unusual dip in PE0.50 in period 20.[17]

Table 6 reports two LPMs using data from the PE0.50 and PE0.80 treatments. Regression (2) is estimated with multi-level random effects at the session-buyer-seller levels.[18] The dependent variable=1 if a contract is binding. Under Empirical Implication 3, the coefficient for the dummy variable for PE0.50 (PE0.80 is the omitted category) should be positive since a drop in δ from 0.80 to 0.50 should increase the probability that the buyer offers a formal contract. Indeed, the coefficients for PE0.50 are positive and significant in both regressions suggesting a robust positive effect.[19]

Next, we examine contracted quality, Q, for trades that remained under relational contracts in the PE treatments. Under Empirical Implication 3, we

[17]Only one of the three PE0.50 sessions lasted as long as 20 periods so the data from period 20 and beyond is from one session and may be more volatile.

[18]Theoretically, there should be no correlation between unobserved buyer heterogeneity and the other independent variables so that random-effects can be used.

[19]We also ran probit regressions and the qualitative results are unchanged, so we do not report them.

Table 6: LPM Estimates (dep. var.=1 if binding contract)

	(1)	(2)
PE50 (dummy)	0.197***	0.174***
	(0.07)	(0.06)
Period	0.041**	0.046***
	(0.017)	(0.017)
Period^2	-0.0004	-0.0008
	(0.0006)	(0.0006)
Constant	-0.04	-0.037
	(0.096)	(0.09)
Random-Effects	No	Session-Buyer-Seller levels
Observations	672	672

-Estimated using data from all sessions for treatments PE0.80 and PE0.50
-Robust standard errors adjusted for clustering on sessions in parentheses
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

expect Q to be higher in PE0.80 vs. PE0.50. However, Figure 3 shows that the average Q is similar across the two treatment ($Q = 9.78$ in PE0.50 versus $Q = 9.95$ in PE0.80).[20]

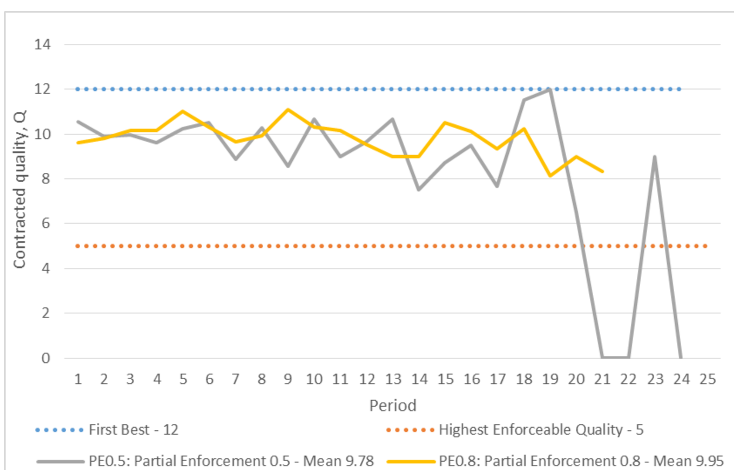

Figure 3: *Average level of Q.*

[20]The volatility of the PE0.50 line reflects the fact that there were very few observations involving relational contracting after period 19 in the PE0.50 sessions.

Table 7 reports regression results isolating the impact of δ on Q. Theoretically, the coefficient for the PE0.50 dummy should be negative. While the estimated coefficient in regression (1) is -0.231, it is not significant. Regression (2) is estimated with random-effects at the session-buyer-seller levels, but the coefficient is still insignificant. Thus, there is no evidence that a reduction in δ from 0.80 to 0.50 affects Q for those subjects who use relational contracts.

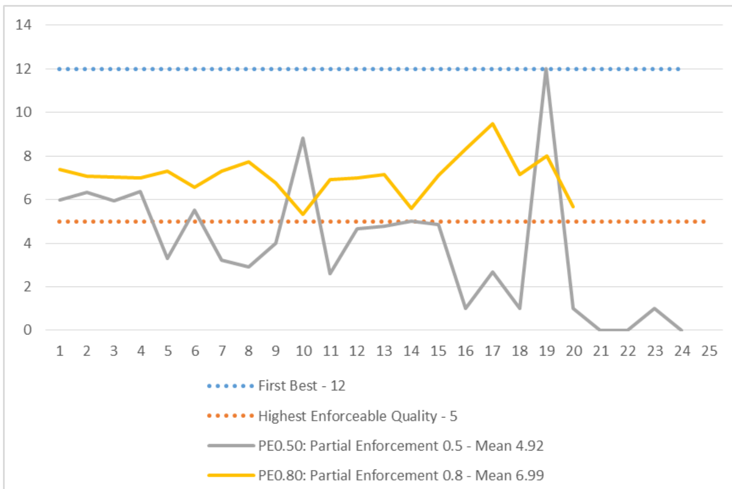

Figure 4: *Average level of q across PE treatments.*

It seems odd that the average Q in PE0.50 is 9.78, which exceeds the maximum self-enforcing level of $Q = 8$. However, in Figure 4, we see that actual quality delivered in PE0.50 is only $q = 4.92$, roughly half of $Q = 9.78$. Moreover, mean q in PE0.50 trends downward over time. When we examine Figure 2 and Figure 4 in combination, the trend is for subjects in PE0.50 to switch to formal contracts over time and for actual q to trend downward for those who continue to use relational contracts. In contrast, we see no downward trend for actual q in PE0.80, although the mean quality supplied, $q = 6.99$, is still lower than the mean quality demanded, $Q = 9.95$.

Regressions (3) and (4) in Table 7 estimate the impact of the PE0.50 dummy on q. Across both regressions, there is a reduction of more than 2 units in q in PE0.50 relative to PE0.80. Thus, while buyers in both treatments specified similar levels of contracted Q, actual q delivered by sellers is

Table 7: Impact of a Drop in δ on Contracted and Actual Quality

	Dependent variables			
	Q	Q	q	q
	(1)	(2)	(3)	(4)
PE0.50	-0.231	-0.332	-2.206***	-2.111***
	(0.154)	(0.323)	(0.442)	(0.474)
Period	0.037	-0.035	-0.192***	-0.277***
	(0.078)	(0.092)	(0.0667)	(0.064)
Period^2	-0.005	-0.002	0.0062	0.01***
	(0.004)	(0.0044)	(0.0049)	(0.0034)
Constant	10.15***	10.40***	8.002***	7.65***
	(0.289)	(0.300)	(0.332)	(0.353)
Random-Effects	No	Session-Buyer-Seller	No	Session-Buyer-Seller
Observations	382	382	382	382

- The omitted category is PE0.80.

-Robust standard errors adjusted for clustering on sessions in parentheses

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

substantially lower in PE0.50. This is not all that surprising considering that the average Q value of 9.78 is theoretically not self-enforcing with $\delta = 0.50$.

An interesting puzzle is why buyers under-specify Q in PE0.80 and yet over-specify Q in PE0.50. Recall that, theoretically, it should be possible to self-enforce $Q = 12$ in PE0.80 and yet buyers specified only $Q = 9.95$, on average. In contrast, it should be possible to self-enforce only a maximum Q of 8 in PE0.50 and yet buyers specified 9.78. We offer a couple of possible explanations and leave a more detailed analysis for future work.

First, because self-enforcement is so difficult in PE0.50, buyers may strategically design contracts for opportunistic purposes with no intention of self-enforcement. This conjecture is supported by the fact that buyers shirked 88% of the time in PE0.50 and sellers shirked 80% of the time. By the later periods, about 60% to 80% of the trades in PE0.50 were conducted with binding contracts so the few that used non-binding contracts may have been experimenting with ways to extract profit in an opportunistic way. One way of

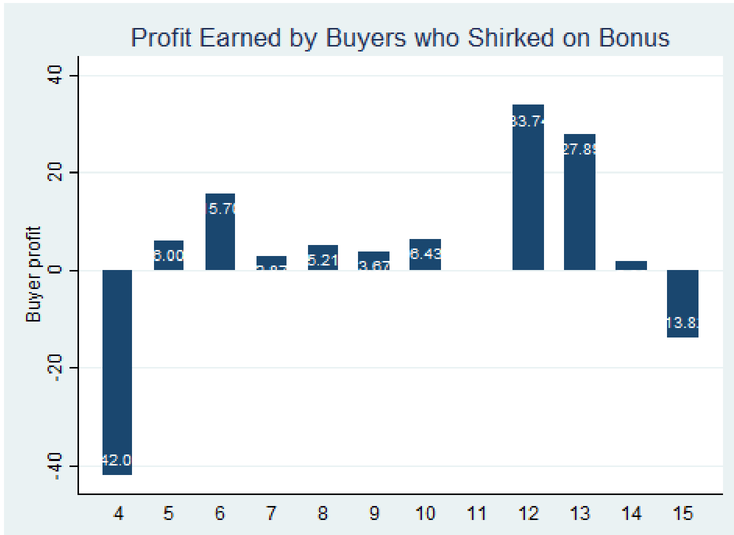

Figure 5: *Profit earned by shirking buyers in PE0.50 across Q*

engaging in opportunism is to ask the seller to deliver a very high level of quality even if the buyer has no intention of honoring the promised bonus. Figure 5 shows that the most profitable opportunistic buyers requested Q in the neighborhood of the first-best value ($Q = 12$).

Second, because self-enforcement is achievable in PE0.80 for Q up to the first-best level, perhaps the main goal of buyers was not to engage in opportunism but to protect against strategic uncertainty. Recall that Breitmoser (2015) suggests that semi-grim strategies do not rule out conflict even after mutual cooperation. In this case, it is natural to choose a lower Q which provides more slack in the SE constraints to ensure mutual performance.

For comparison, we can use the data from treatment E to examine behavior in the absence of strategic uncertainty. A key characteristic of treatment E is that buyers can use formal contracts to implement any quality level without fear of strategic uncertainty because the computer ensures that $Q = q$. Figure 6 shows that binding contracts in treatment E implemented mean actual quality remarkably close to the first best level of 12 (11.7 vs 12). Moreover, 48% of trades resulted in exactly the first-best quality. The few incomplete contracts used implemented $q = 7.14$ with only 5% implementing the first best.[21]

[21] Moreover, the incomplete contracts plot was volatile because very few trades used in-

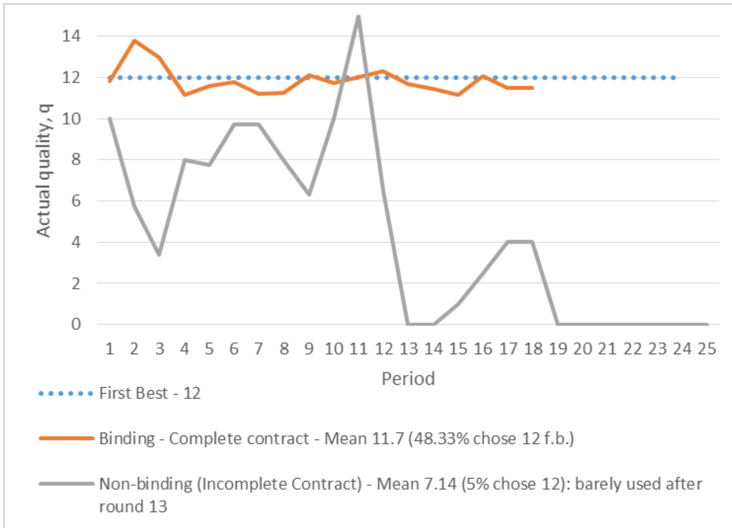

Figure 6: *Actual q Realized in Treatment E - Perfect 3rd Party Enforcement*

Thus, when strategic uncertainty is eliminated, subjects chose Q that were remarkably close to the first best even though the first-best value of 12 was an interior solution and not an obvious focal point.

4.4 Empirical Implication 4: Impact of Verifiability

Empirical Implication 4 predicts that subjects will use relational contracts rather than formal contracts when verifiability is imperfect; i.e., more complete contracts should be observed in the E relative to the PE treatments.

Recall that the descriptive statistics in Figure 2 are consistent with Empirical Implication 4 in that the E–treatment yields a significantly higher fraction of complete contracts than either PE0.50 or PE0.80. The means are 0.81, 0.54 and 0.35, respectively, and the gaps appear to persist across almost all periods. In short, with partial enforcement, a large number of contract offers leave out even costlessly verifiable terms such as contractible quality $Q = 5$.

While Figure 2 provides an overview of the results, we also conduct formal hypothesis testing. Table 8 contains the results from two LPMs. Regression

complete contracts. In many periods, only one or two trades were executed using incomplete contracts. In the later periods, many trades did not use incomplete contracts at all. These are the observations for which the plot touched zero quality.

Table 8: LPM Estimates (dep. var.=1 if binding contract)

	(1)	(2)
PE	-0.380***	-0.322***
	(0.067)	(0.086)
1-memory cooperation dummy		-0.005
		(0.045)
Period	0.036**	0.076***
	(0.0139)	(0.014)
$Period^2$	-0.0002	-0.002 ***
	(0.0005)	(0.0005)
Constant	0.465***	0.296***
	(0.079)	(0.079)
R^2	0.25	-
Observations	893	551

-Estimated using data from all sessions for all treatments (E, PE0.80 and PE0.50)
-Robust standard errors adjusted for clustering on sessions in parentheses
-Regression (2) is a multi-level random-effects linear probability model at the session-buyer-seller levels.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

(2) adds the 1-memory cooperation dummy and is estimated with multi-level random-effects at the session-buyer-seller levels.[22] The dependent variable takes a value of 1 if a binding complete contract is offered. Both regressions include the treatment dummy for PE (either from PE50 or PE80). The base category is the E treatment and therefore the sign of the coefficient for PE tests Empirical Implication 4. A negative sign is expected since it suggests that the probability of a complete contract offer decreases under partial enforcement.

The estimated coefficients for the PE dummy are both negative and significant. Moreover, the magnitude of the estimated coefficients are similar suggesting robustness. Thus, we cannot reject the prediction that people move

[22]We also separately estimated fixed effects and probits, but the qualitative results were unchanged.

toward relational contracts when facing verifiability imperfections.

4.5 Empirical Implication 5: Contractual Form

Empirical Implication 5 predicts that we should observe discretionary bonus contracts rather than efficiency wage/gift-exchange type fixed-price contracts when subjects can endogenously choose contractual form.

Figure 7 shows that subjects overwhelmingly traded using discretionary bonus contracts rather than efficiency wage/fixed price contracts. Despite the popularity of fixed price contracts in the experimental economics literature, it appears that, when given a choice, few subjects use these types of contracts. This is consistent with the insights of Bernheim and Whinston (1998).

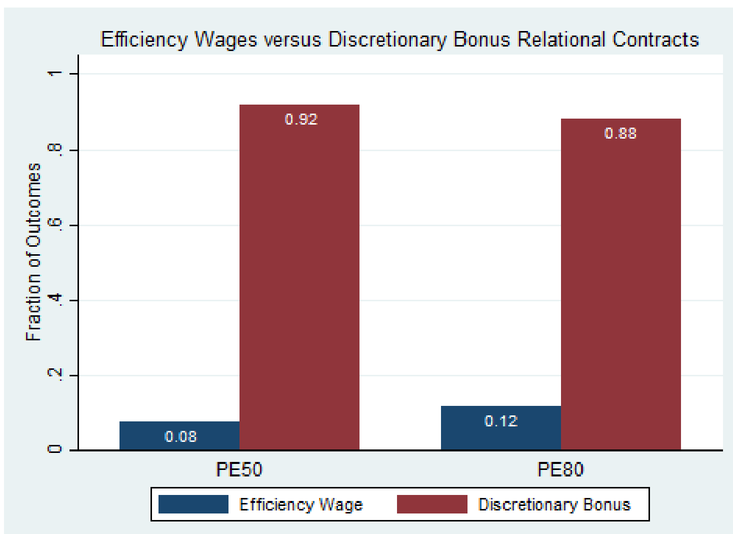

Figure 7: *Use of efficiency wage versus discretionary bonus contracts.*

Figure 8 shows the profits buyers earn from offering efficiency wage versus discretionary bonus contracts. Buyers earned substantially higher profit under discretionary bonus contracts across both treatments. Thus, it is not surprising that buyers preferred discretionary bonus contracts.

Next, we examine efficiency under both types of contracts. Table 9 reports regression results to examine the impact of discretionary bonus contracts on actual quality outcomes, q. The omitted category is the efficiency wage. Re-

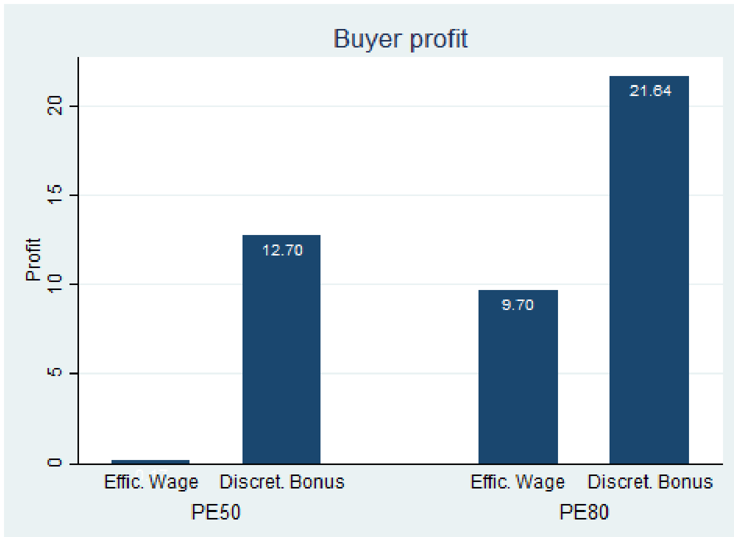

Figure 8: *Buyer profit under efficiency vs discretionary bonus contracts.*

Table 9: Impact of discretionary bonus contracts on quality (q)

	(1)	(2)
Discretionary bonus contract dummy	2.07*	4.42*
	(0.86)	(1.82)
PE0.80 dummy	2.33***	10.52***
	(0.471)	(1.86)
Lagged cooperation dummy		2.13***
		(0.365)
Period	-0.212***	-0.141
	(0.051)	(0.097)
Period^2	0.007	0.006
	(0.005)	(0.007)
Constant	3.94***	-0.16
	(0.73)	(1.73)
Seller fixed effects	No	Yes
Observations	382	189

-Robust standard errors adjusted for clustering on sessions in parentheses
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

gression (2) differs from (1) in that it includes seller fixed effects and the 1-memory cooperation dummy to control for possible seller selection issues and belief updating about buyer actions.

We point out that first best is $q^* = 12$ and average quality across both contract types in both treatments fell short of 12. Thus, when interpreting the coefficients, a positive coefficient for the discretionary bonus dummy would imply that the discretionary bonus contract induces more efficiency via higher quality. Both regressions show that discretionary bonus contracts increases quality (2.07 in regression (1) and 4.42 in regression (2)) which are statistically significant at the 10% level. Thus, the regressions provide more evidence about why buyers tend to choose discretionary bonus contracts.

Overall, our results support Empirical Implication 5 and provide empirical justification for the optimality of discretionary bonus contracts and the importance of strategic flexibility in incomplete contracts in the sense of BW.

5 Conclusion

We use economic experiments to test a number of well-known empirical implications from canonical relational contract theory. Our results support the majority of the implications and suggest that standard relational contracting theory is useful for explaining many empirical patterns of behavior. When total pay does not meet individual rationality conditions or the promised discretionary bonus does not satisfy the agent's incentive compatibility condition, there is an increase in contract rejection or shirking. With partial enforcement, subjects rely more on relational contracts. With a drop in discount factor, subjects shift to formal contracts in the partial enforcement treatments. Finally, in the presence of imperfect verifiability, subjects largely choose discretionary bonus contracts rather than efficiency wage contracts, which is consistent with the theoretical optimality of discretionary bonus contracts in our model and the theory of strategic ambiguity of Bernheim and Whinston (1998).

Despite the success of the standard theory in explaining most patterns of behavior in our experiments, we also came across some surprising results. We provide, to our knowledge, the first empirical investigation of parties' post-shirking strategies within the context of relational contracts. Here, our results do not support the strict prediction that principals only use efficient punishments (in the sense of the "strongly optimal" contracts of Levin, 2003) following deviations. Nonetheless, our results are consistent with a "less grim" version of the semi-grim strategies of Breitmoser (2015), which relies on belief-free equilibria to explain different probabilities of cooperation given different 1-memory histories.

We also find that, contrary to theoretical predictions, subjects assigned to the role of agents are reluctant to reject contracts with excessively large bonuses that breach the principal's self-enforcement constraint. Similarly, subjects assigned to the role of principals do not react optimally to a decrease in the discount factor by demanding a lower quality in relational contracts. These differences between our findings and theoretical predictions offer opportunities to extend theory by integrating behavioral insights or insights from recent developments in the theory of repeated games on semi-grim strategies.

References

Aghion, P., and R. Holden. 2011. "Incomplete Contracts and the Theory of the Firm: What Have We Learned Over the Past 25 Years?" *The Journal of Economic Perspectives*, 25: 181–197.

Akerlof, G.A. 1982. "Labor contracts as partial gift exchange." *The Quarterly Journal of Economics*, 97(4): 543–569.

Antras, P., and F. Foley. 2015. "Poultry in motion: a study of international trade finance practices." *Journal of Political Economy (forthcoming)*.

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

Bernheim, B.D., and M.D. Whinston. 1998. "Incomplete contracts and strategic ambiguity." *American Economic Review*, 902–932.

Boot, A.W.A., S.I. Greenbaum, and A.V. Thakor. 1993. "Reputation and discretion in financial contracting." *American Economic Review*, 1165–1183.

Bó, Pedro Dal. 2005. "Cooperation under the shadow of the future: experimental evidence from infinitely repeated games." *American Economic Review*, 1591–1604.

Breitmoser, Yves. 2015. "Cooperation, but no reciprocity: individual strategies in the repeated prisoner's dilemma." *American Economic Review*, 105(9): 2882–2910.

Brown, M., A. Falk, and E. Fehr. 2004. "Relational contracts and the nature of market interactions." *Econometrica*, 72(3): 747–780.

Brown, Martin, Armin Falk, and Ernst Fehr. 2012. "Competition and relational contracts: the role of unemployment as a disciplinary device." *Journal of the European Economic Association*, 10(4): 887–907.

Casoria, F., and A. Riedl. 2013. "Experimental labor markets and policy considerations: Incomplete contracts and macroeconomic aspects." *Journal of Economic Surveys*, 27(3): 398–420.

Fehr, E., A. Klein, and K.M. Schmidt. 2007. "Fairness and contract design." *Econometrica*, 75(1): 121–154.

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

Gächter, Simon, and Armin Falk. 2002. "Reputation and reciprocity: Consequences for the labour relation." *The Scandinavian Journal of Economics*, 104(1): 1–26.

Gil, R., and G. Zanarone. 2015. "On the determinants and consequences of informal contracting." *Available at SSRN 2560520.*

Halac, M. 2012. "Relational Contracts and the Value of Relationships." *American Economic Review*, 102: 750–779.

Klein, B., and K.B. Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance." *The Journal of Political Economy*, 615–641.

Levin, J. 2003. "Relational incentive contracts." *The American Economic Review*, 93(3): 835–857.

Macchiavello, R., and A. Morjaria. 2015. "The value of relationships: evidence from a supply shock to Kenyan Rose Exports." *American Economic Review*, 105(9): 2911–2945.

MacLeod, W.B. 2007. "Reputations, Relationships, and Contract Enforcement." *Journal of Economic Literature*, 595–628.

MacLeod, W.B., and J.M. Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica*, 447–480.

Schmitz, K.M., and M. Schnitzer. 1995. "The interaction of explicit and implicit contracts." *Economic Letters*, 48: 193–199.

Scott, R.E. 2003. "Theory of Self-Enforcing Indefinite Agreements." *Colum. Law Review*, 103: 1641.

Shapiro, C., and J.E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review*, 74: 433–444.

Sloof, R., and J. Sonnemans. 2011. "The interaction between explicit and relational incentives: An experiment." *Games and Economic Behavior*, 73: 573–594.

Telser, L.G. 1980. "A Theory of Self-Enforcing Agreements." *Journal of Business*, 53: 27–44.

Yang, H. 2013. “Nonstationary Relational Contracts with Adverse Selection.” *International Economic Review*, 54(2): 525–547.

APPENDIX - FOR ONLINE PUBLICATION

A Full Description of the Theoretical Model, Propositions, and Implications

A.1 Model setup

We describe a simple model that can conceptualize many of the standard predictions from the relational contracting literature. Our purpose is not to derive new theory but rather to provide a parsimonious unifying framework for many canonical results that span several papers in the literature. Such a simple, unifying framework serves the dual purpose of providing clear intuitive predictions and facilitating laboratory implementation where simplicity is not only a virtue, but a necessity.

Assume that a principal contracts with an agent to produce a unit of a good for which quality matters. For simplicity, we abstract from asymmetric information, so our environment is similar to MM where the key friction is the absence of third-party enforcement.

The agent's obligation is to deliver quality $q \geq Q$ where Q refers to the quality level specified in the contract and q refers to the actual quality delivered. The principal's obligation is to pay $w \geq W$ where w is actual payment and W is the payment specified in the contract. w can consist of a base price p and bonus payment b, so we write $w = p + b$. Similarly, we write $W = P + B$ for the contractually specified payments. Since P is a fixed, non-contingent payment, p=P by default.

Let the principal's and agent's payoffs be $\pi_P = r(q) - p - b$ and $\pi_A = p + b - c(q)$ where $r(q)$ and $c(q)$ are differentiable functions such that $r'(q) > 0$, $r''(q) \leq 0$, $c'(q) > 0$ and $c''(q) \geq 0$, $\forall q \in [\underline{q}, \overline{q}] \subset \mathbb{R}_+$. All else equal, the principal prefers higher quality and lower payments, and the agent prefers higher payment and lower quality. The reservation payoffs for the principal and agent are $\overline{\pi}$ and $\overline{u}$, respectively. Assume that there exists some minimal quality threshold $\check{q} \in (\underline{q}, \overline{q})$ such that $r(q^h) - c(q^h) \geq \overline{u} + \overline{\pi} > r(q^l) - c(q^l)$ for

$q^l \in [\underline{q}, \check{q})$ and $q^h \in [\check{q}, \overline{q}]$. This implies a minimum quality must be produced to generate positive surplus.

A.2 Nesting Formal and Relational Contracts

We assume limited third-party verifiability where a third-party is able to detect whether the good achieves some coarse, discrete level of quality but cannot detect more refined gradations in quality. Limited third-party verifiability allows for imperfections in performance measurement in the spirit of BGM, but it conceptualizes the issue in a simpler one-dimensional framework that facilitates experimental implementation. Moreover, in practice, many products receive discrete quality certifications that are neither completely unenforceable by a third-party nor enforced to highly refined quality grades. Thus, our setup better matches stylized observations while allowing for a nesting of both formal and informal contracts in a parsimonious framework.

Enforcement imperfections do not preclude the possibility of writing formal/complete contracts, though imperfections do limit the set of available complete contracts. Partition the quality space $[\underline{q}, \overline{q}] \in \mathbb{R}_+$ into $[[\underline{q}, q^d), [q^d, \overline{q}]]$ where q^d is a quality threshold that can be feasibly verified by a third-party.

Assumption 1. *A third-party can verify whether* $q \in [\underline{q}, q^d)$ *or* $q \in [q^d, \overline{q})$

Assumption 1 implies a contractible set, $\underline{C} = \{\underline{q}, q^d\}$. No other quality level is verifiable; hence, the agent will choose $q = q^d$ even if a contract calls for $Q > q^d$ and will choose $q = \underline{q}$ if the contract calls for $\underline{q} < Q < q^d$.

Despite imperfect enforcement, it is still possible to write a formal contract. A formal contract must be a complete contract in that a complete state-contingent plan governs performance. Therefore, all obligations of both parties are fully specified for all contingencies in the initial contract. Moreover, the contract is third-party enforceable so that neither party can shirk. This implies that no party has ex post discretionary latitude to deviate from the initial contract. One can view the presence of ex post discretion to deviate as being synonymous with an incomplete contract. This implies that the contract would have to be self-enforcing through an informal agreement.

The complete contract can either specify state-contingent prices $\underline{P}$ and P^d to be paid under each contractible quality realization, or the principal can specify $Q = q^d$ in exchange for a fixed P. We will refer to the latter as a **simple contract**. In the former case, a third-party enforces the contingent payments $\underline{P}$ and P^d whereas in the simple contract, $Q = q^d$ and P are directly enforced. In either case, all variables are third-party enforceable since they are either in the contractible set or depend only on variables in the contractible set. If the contingent payments $\underline{P}$ and P^d are chosen in an incentive compatible manner to implement $Q = q^d$, then the two types of contracts are outcome equivalent. Thus, for simplicity, we will focus only on simple contracts.

We also describe incomplete contracts to frame our subsequent discussion of optimal relational contracts and strategic incompleteness. Note that there is no unique incomplete contract, so we illustrate one example. Suppose a contract specifies $Q > q^d$, a fixed payment P and a bonus B if $q \geq Q$ is realized. Because $Q > q^d$ is not in the contractible set, it follows that the agent has ex post discretion to deviate from Q without legal consequence. Additionally, because B is contingent on $q \geq Q$, B is a discretionary bonus that is not contractible. Therefore, the principal can shirk on the bonus even if the agent performs. In summary, both parties have ex post discretion to deviate from the initial agreement. Backward induction shows that our illustrated incomplete contract above leads to inefficiencies in the absence of self-enforcement.

To model endogenous incompleteness, we denote π^f and u^f as the payoffs obtained from the "best" complete contract for the given enforcement technology; i.e., the formal contract that yields the highest joint surplus under the enforcement technology. In our case, if the first best quality level is such that $q^* > q^d$, then a formal contract specifying q^d would dominate one specifying $\underline{q}$. Since there are only two contractable quality levels, the contract specifying q^d is the best complete contract. Denote Q^f as the best contracted quality level.[23] Denote surplus as $S(q) = r(q) - c(q) - \overline{u} - \overline{\pi}$. We define

$$k = S(q^*) - S(Q^f) \tag{A.1}$$

[23] In our example $Q^f = q^d$.

to be the loss in efficiency from using a formal contract in the presence of verifiability imperfections. Note that when a third-party can verifiy every quality level, then $k = 0$ since $Q^f = q^*$.

Like BGM, our model nests formal and informal contracts. Unlike BGM, we have a single performance measure rather than separately defining objective and subjective measures. This setup eases experimental implementation since subjects track fewer variables.

A.3 Optimal Contracting

Consider a principal-agent model of repeat trading under the imperfect enforcement technology specified above. We define a binary variable $\alpha \in \{0, 1\}$ where α equals 1 if $u^f + \pi^f \geq \overline{u} + \overline{\pi}$ and 0 otherwise. That is, $\alpha = 1$ if joint profits from the best complete contract exceeds joint reservation payoffs. The stage-game timeline follows the typical principal-agent sequence:

1. Principal offers a contract–a price/bonus/quality triplicate, (P, B, Q).
2. The agent accepts or rejects. If rejected, the parties default to the best formal contract if $\alpha = 1$ and to reservation payoffs if $\alpha = 0$.
3. If accepted, the agent chooses actual quality q.
4. The principal observes q and chooses actual bonus b. The promised fixed payment, P, is also made.[24]

A relational contract is an infinite repetition of the above stage-game so that in each period t and for each history up to t, the relational contract describes the sequence (1)-(4). Moreover, the relational contract is self-enforcing if it describes a subgame perfect equilibrium of the infinitely repeated game. In addition, Levin (2003) and Halac (2012) show that, with symmetric information, there exist stationary contracts that are optimal in that the same

[24] P is always third party enforceable because it is not contingent on quality.

(optimal) contract is offered in every t.[25] Letting δ be the discount factor and multiplying the payoffs by $1-\delta$ to express them as per-period averages, the principal's contract design problem is:

$$\max_{Q,P,B}(1-\delta)\left[r(Q)-P-B\right]+\delta V\left[C\right] \quad s.t. \tag{A.2}$$

$$(1-\delta)\left[r(Q)-P-B\right]+\delta V\left[C\right]\geq \alpha\pi^{f}+(1-\alpha)\overline{\pi} \tag{A.3}$$

$$(1-\delta)\left[P+B-c(Q)\right]+\delta U\left[C\right]\geq \alpha u^{f}+(1-\alpha)\overline{u} \tag{A.4}$$

$$(1-\delta)\left[r(Q)-P-B\right]+\delta V\left[C\right]\geq (1-\delta)\left[r(Q)-P\right]+\delta\left[\alpha\pi^{f}+(1-\alpha)\overline{\pi}\right] \tag{A.5}$$

$$(1-\delta)\left[P+B-c(Q)\right]+\delta U\left[C\right]\geq (1-\delta)\left[P-c(\underline{q})\right]+\delta\left[\alpha u^{f}+(1-\alpha)\overline{u}\right] \tag{A.6}$$

Constraints A.3 and A.4 are the individual rationality (IR) constraints and A.5 and A.6 are the self-enforcement (SE) constraints. To understand the expressions $V(C)$ and $U(C)$, let Γ denote the set of feasible contracts, which can be partitioned as $C\cup F=\Gamma$ and $C\cap F=\emptyset$. Then, either $(P,B,Q)\in C$ or F, where "C" denotes relational contracts that satisfy contraints A.3-A.6, and "F" denotes "formal" (i.e., complete) contracts that only satisfy the IR constraints. Thus, $V(C)$ and $U(C)$ are the flow payoffs for the principal and agent, respectively, from the optimal self-enforcing relational contract $(P,B,Q)\in C$. Due to stationarity, the same contract is offered every t, so the principal's contract design problem becomes essentially a static optimization problem.

Proposition 1. *Solving problem A.2 yields an optimal stationary contract that requests $\tilde{Q}\leq Q^{*}$ where Q^{*} is a request for first best quality. The associated*

[25] Nonstationary contracts arise primarily in the context of private information where one has to model relational dynamics due to the revelation of private information over time (e.g., see Halac, 2012 or Yang, 2013). It's important to point out that nearly all experiments involve some dynamics simply because subjects learn how to play the game. Hence, researchers typically treat predictions from stationary symmetric information games as theoretical benchmarks that subjects should converge to after sufficient learning. The actual dynamics that lead to convergence is typically not of theoretical interest and early period departures from theoretical benchmarks are treated as noise that can be reduced with subject experience.

payment scheme is $W(\tilde{Q}) = \tilde{P} + B(\tilde{Q})$ *such that:*

(i) $\frac{\alpha u^f+(1-\alpha)\overline{u}+c(\tilde{Q})}{1-\delta} - \frac{\delta}{1-\delta}\{r(\tilde{Q}) - \alpha\pi^f - (1-\alpha)\overline{\pi}\} \leq \tilde{P} \leq \alpha u^f + (1-\alpha)\overline{u} + c(\underline{q})$

(ii) $c(\tilde{Q}) - c(\underline{q}) \leq B(\tilde{Q}) \leq \frac{\delta}{1-\delta}\{r(\tilde{Q}) - c(\tilde{Q}) - \alpha\pi^f - (1-\alpha)\overline{\pi} - \alpha u^f - (1-\alpha)\overline{u}\}$

(iii) $\tilde{P} + B(\tilde{Q}) - c(\tilde{Q}) \geq \alpha u^f + (1-\alpha)\overline{u}$

(iv) $r(\tilde{Q}) - \tilde{P} - B(\tilde{Q}) \geq \alpha\pi^f + (1-\alpha)\overline{\pi}$

Proof. First note that with stationary contracts, this essentially becomes a static problem since $V(C) = r(Q) - P - B$ at the optimal self-enforcing values of (Q, P, B). Second, note that A.5 and A.6 can be combined to get:

$$\frac{\delta}{1-\delta}\left[V(C) - \alpha\pi^f - (1-\alpha)\overline{\pi}\right] \geq B \geq \left[c(Q) - c(\underline{q})\right] - \frac{\delta}{1-\delta}\left[U(C) - \alpha u^f - (1-\alpha)\overline{u}\right] \tag{A.7}$$

Additionally, A.7 can be rearranged to get:

$$\frac{\delta}{1-\delta}[r(Q) - c(Q) - \alpha\pi^f - (1-\alpha)\overline{\pi} - \alpha u^f - (1-\alpha)\overline{u}] \geq c(Q) - c(\underline{q}) \tag{A.8}$$

Given the assumptions $r'(Q) \geq 0, r''(Q) \leq 0, c'(Q) > 0$, and $c''(Q) \geq 0$, A.8 tightens as Q increases. Suppose that $\hat{Q}$ is the value of Q at which A.8 holds with equality. Then if $Q^* > \hat{Q}$, then Q^* is not implementable. However, if $Q^* \leq \hat{Q}$, then Q^* can be implemented. Therefore, the principal can only contract for some $\tilde{Q} \leq Q^*$.

To derive the optimal payment scheme, we must consider two cases. First, if $\hat{Q} \geq Q^*$ so that the principal can contract for the first best level of quality where $r'(Q^*) = c'(Q^*)$, then there is slack in A.7. Second, if $\hat{Q} < Q^*$ so $r'(\hat{Q}) > c'(\hat{Q})$, then the principal will contract for $\tilde{Q} = \hat{Q}$ and A.7 binds with equality. We will analyze each case separately.

Case 1: $\hat{Q} \geq Q^*$: In this case, there is slack in A.7 even when $\tilde{Q} = Q^*$. To maintain self-enforcement, the principal can offer any $B(\tilde{Q})$ in the interval $\frac{\delta}{1-\delta}\left[V(C) - \alpha\pi^f - (1-\alpha)\overline{\pi}\right] \geq B(\tilde{Q}) \geq \left[c(\tilde{Q}) - c(\underline{q})\right] - \frac{\delta}{1-\delta}\left[U(C) - \alpha u^f - (1-\alpha)\overline{u}\right]$. This is consistent with (ii). Moreover, P must be chosen in combination with $B(\tilde{Q})$ to obey both the principal's and agent's individual rationality constraints. This is consistent with (iii) and (iv).

Case 2: $\hat{Q} < Q^*$: Then $r'(\hat{Q}) > c'(\hat{Q})$ so the maximum self-enforcing $\tilde{Q}$

that the principal can contract for is $\hat{Q}$. The corresponding self-enforceable $B(\tilde{Q}) = \frac{\delta}{1-\delta}[r(\tilde{Q}) - c(\tilde{Q}) - \alpha\pi^f - (1-\alpha)\overline{\pi} - \alpha u^f - (1-\alpha)\overline{u}] = c(\tilde{Q}) - c(\underline{q})$, which satisfies part (ii) with equality. P must be chosen in combination with $B(\tilde{Q})$ to obey both the principal's and agent's individual rationality constraints. This is consistent with (iii) and (iv).

□

In words, under the optimal contract, the principal contracts for quality that is less than or equal to first best quality; the discretionary bonus simultaneously satisfies both the agent's and principal's SE constraints; and the total promised payment satisfies both parties' IR constraints. This leads directly to Empirical Implication 1 in the main body of the paper.

For a more intuitive look at self-enforcement, we can also solve the expression in Proposition 1(ii) for δ which yields:

$$\delta \geq \underline{\delta}(Q) = \frac{c(Q) - c(\underline{q})}{r(Q) - c(\underline{q}) - \alpha\left[\pi^f + u^f\right] - (1-\alpha)\left[\overline{\pi} + \overline{u}\right]} \tag{A.9}$$

$$= \frac{c(Q) - c(\underline{q})}{r(Q) - c(\underline{q}) - \alpha\left[r(Q^f) - c(Q^f)\right] - (1-\alpha)\left[\overline{\pi} + \overline{u}\right]} \tag{A.10}$$

$\underline{\delta}(Q)$ is the threshold for the incomplete contract to be self-enforcing, and it depends on Q, where a higher Q raises the threshold making self-enforcement more difficult. Consequently, this can limit the quality that can be implemented. The threshold also depends on the payoffs u^f and π^f, which in turn, depends on the efficiency loss from imperfect verifiability. Thus, self-enforcement and third-party enforcement interact; i.e. suppose Q^f is the enforceable quality that yields the highest joint surplus among all contractible quality levels. A complete contract (Q^f, P^f) yields payoffs $\pi^f = P^f - c(Q^f)$ and $u^f = P^f - c(Q^f)$. These payoffs can be substituted in (A.9) to get (A.10). As k in (A.1) tends toward zero, third-party verifiability improves. This, in turn, increases the joint profit $r(Q^f) - c(Q^f)$ which weakly raises the threshold for self-enforcement A.9.[26] In short, an improvement in enforcement technology should cause some relational contracts to be replaced by complete contracts.

Proposition 2. *Let Q^* be the first best quality request such that $Q^* \in \arg\max_Q \{S(Q)\}$. If there exists $\tilde{Q}$ such that $S(Q^*) \geq S(\tilde{Q}) > max\{S(Q^f), \overline{\pi} + \overline{u}\}$ and $\delta \geq \underline{\delta}(\tilde{Q})$, then a relational contract that implements $\tilde{Q}$ is preferred over the best complete contract or termination.*

Proof. If there exists $\tilde{Q}$ such that $S(Q^*) \geq S(\tilde{Q}) > S(Q^f)$ and $\delta \geq \underline{\delta}(\tilde{Q})$, then $\tilde{Q}$ is a self-enforcing level of quality that yields higher surplus than the best complete contract. Thus, the principal can allocate enough surplus to both parties to make them at least as well off as they would be under the best complete contract. Hence, $\tilde{Q}$ is a self-enforcing quality level that satisfies constraints A.3-A.6 and can be made jointly preferred by the principal and agent. □

Proposition 2 states that if verifiability is sufficiently imperfect, which allows for the existence of some self-enforcing level of $\tilde{Q}$ that yields joint surplus that is greater than joint surplus under the other options, then the parties will use relational contracts.

Levin (2003)'s Corollary 1 (p. 841) points out that, because optimal stationary contracts can be constructed to split the surplus in any way the parties desire (subject to IR constraints), the parties can continue with a relational contract even following a deviation.

Corollary 1. *Following any history, there exists a family of optimal relational contracts that implements $\tilde{Q}$ such that $S(\tilde{Q}) > max\{S(Q^f), \overline{\pi} + \overline{u}\}$ and yield per-period payoffs $\tilde{\pi} \in [max\{\pi^f, \overline{\pi}\}, S(\tilde{Q}) - max\{u^f, \overline{u}\}] \subset \mathbb{R}$ to the principal, and per-period payoffs $\tilde{u} = S(\tilde{Q}) - \tilde{\pi}$ to the agent.*

Proof. Any contract that implements $\tilde{Q}$ and yields per-period payoffs $\tilde{\pi} \in [max\{\pi^f, \overline{\pi}\}, S(\tilde{Q}) - max\{u^f, \overline{u}\}]$ to the principal, and per-period payoffs $\tilde{u} =$

[26] We say weakly because if $\alpha = 0$, then the threshold does not change until complete contracts joint surplus exceeds joint surplus from the reservation payoffs, triggering $\alpha = 1$.

$S(\tilde{Q}) - \tilde{\pi}$ to the agent satisfies all the conditions enumerated in Proposition 1 and is therefore optimal. Moreover, by Proposition 2, $S(\tilde{Q}) > max\{S(Q^f), \overline{\pi} + \overline{u}\}$. Thus, for any history in which both parties honor this contract ($q \geq \tilde{Q}$ and $b \geq B(\tilde{Q})$), the parties continue with this contract by stationarity.

For any history in which at least one party deviates ($q < \tilde{Q}$ and/or $b < B(\tilde{Q})$), there is no need to resort to termination or a formal contract because an optimal relational contract can be constructed by raising P to yield per-period payoffs of $\tilde{\pi} = max\{\pi^f, \overline{\pi}\}$ and $\tilde{u} = S(\tilde{Q}) - max\{\pi^f, \overline{\pi}\}$ if the principal deviated, or by lowering P to yield per-period payoffs $\tilde{\pi} = S(\tilde{Q}) - max\{u^f, \overline{u}\}$ and $\tilde{u} = max\{u^f, \overline{u}\}$ if the agent deviated. Such a contract continues to implement $\tilde{Q}$ because the self-enforcing conditions (part (ii) of Proposition 1) is independent of P. Such a contract provides punishments that are payoff equivalent to termination or reversion to a formal contract. □

Corollary 1 is a modified version of Levin (2003)'s "strongly optimal" contract for our problem. It states that following any history, including those that are off-the-equilibrium path (i.e., a deviation), there is a family of relational contracts that implement $\tilde{Q}$ while delivering different payoff distributions. Thus, one can always construct an off-the-equilibrium path contract that continues to implement $\tilde{Q}$, while holding the deviator to the payoff he would have received had the parties reverted to a formal contract or termination. In other words, the deviator can be punished as severely as termination of the relational contract, but without destroying surplus and without also punishing the non-deviator. Such a contract does not destroy surplus since surplus is higher under $\tilde{Q}$ than under Q^f or termination and is therefore renegotiation proof. In short, continuing with a relational contract is optimal regardless of whether the parties have deviated or not in the previous period. This leads directly to Empirical Implication 2 in the main paper.

Corollary 2. *(Exogenous change in k) Let $\tilde{Q} \in \tilde{\mathbb{Q}} = \{\tilde{Q} : S(Q^*) \geq S(\tilde{Q}) > S(Q^f)\}$. As $k \to 0$, then $\underline{\delta}(\tilde{Q}) \to 1$ for any $\tilde{Q} \in \tilde{\mathbb{Q}}$ and all incomplete contracts are endogenously replaced with complete contracts.*

Proof. First, note that $k = S(Q^*) - S(Q^f) = S(Q^*) - S(Q^f) = r(Q^*) -$

$c(Q^*) - \overline{u} - \overline{\pi} - r(Q^f) + c(Q^f) + \overline{u} + \overline{\pi} = r(Q^*) - c(Q^*) - [r(Q^f) - c(Q^f)]$. Therefore, $k \to 0$ implies that $r(Q^*) - c(Q^*) - [r(Q^f) - c(Q^f)] \to 0$. Moreover, because $r(Q^*) - c(Q^*) - [r(\tilde{Q}) - c(\tilde{Q})] < r(Q^*) - c(Q^*) - [r(Q^f) - c(Q^f)]$ for all $\tilde{Q} \in \tilde{\mathbb{Q}}$, we also have $r(Q^*) - c(Q^*) - [r(\tilde{Q}) - c(\tilde{Q})] \to 0$ and $r(\tilde{Q}) - c(\tilde{Q}) - [r(Q^f) - c(Q^f)] \to 0$ as $k \to 0$. Next, by assumption, $S(Q^*) = r(Q^*) - c(Q^*) - \overline{u} - \overline{\pi} > 0$. Thus, there exists some $\overline{k}$ such that for $k < \overline{k}$, we have $\alpha = 1$ and A.9 becomes $\frac{c(\tilde{Q}) - c(\underline{q})}{r(\tilde{Q}) - c(\underline{q}) - \left[r(Q^f) - c(Q^f)\right]}$. The latter term can be rewritten as $\frac{c(\tilde{Q}) - c(\underline{q})}{r(\tilde{Q}) - c(\tilde{Q}) - \left[r(Q^f) - c(Q^f)\right] + c(\tilde{Q}) - c(\underline{q})} = \frac{c(\tilde{Q}) - c(\underline{q})}{\left[c(\tilde{Q}) - c(\underline{q})\right]\left[\frac{r(\tilde{Q}) - c(\tilde{Q}) - \left[r(Q^f) - c(Q^f)\right]}{c(\tilde{Q}) - c(\underline{q})} + 1\right]} = \frac{1}{\left[\frac{r(\tilde{Q}) - c(\tilde{Q}) - \left[r(Q^f) - c(Q^f)\right]}{c(\tilde{Q}) - c(\underline{q})} + 1\right]}$. Since $r(\tilde{Q}) - c(\tilde{Q}) - [r(Q^f) - c(Q^f)] \to 0$ as $k \to 0$ and the limit of $c(\tilde{Q}) - c(\underline{q})$ is some finite positive number, $\lim_{k \to 0} \underline{\delta}(\tilde{Q}) = \lim_{k \to 0} \frac{1}{\left[\frac{r(\tilde{Q}) - c(\tilde{Q}) - \left[r(Q_f) - c(Q_f)\right]}{c(\tilde{Q}) - c(\underline{q})} + 1\right]} = 1$ □

Corollary 2 is related to the theory of *strategic ambiguity* of BW and to the substitutability between formal and informal contracts of BGM. BW show that in the presence of verifiability imperfections, parties may deliberately eschew formal contracts so that they can achieve better outcomes by using discretionary flexibility to punish and reward non-verifiable performance. Corollary 2 leads to Empirical Implication 4 in the main paper.

Another BW insight is that, given that contracts must be incomplete, it may be optimal for parties to *increase* the degree of incompleteness. Intuitively, under an incomplete contract, the agent has ex post discretionary latitude to shirk. Thus, the principal may also leave herself with discretion via a discretionary bonus contract so that she can adjust pay in response to the agent's action. Such a contract is less complete than a fixed-price contract because the fixed-price contract locks down the principal's obligations. Fixed price contracts are commonly invoked in the literature under the assumption that parties to a relational contract use efficiency wages or repeat purchase mechanisms (Klein and Leffler, 1981; Shapiro and Stiglitz, 1984; Brown, Falk and Fehr, 2004). However, Proposition 1 supports the theory of strategic ambiguity rather than a fixed-price contract. This leads directly to Empirical Implication 5 in the main paper.

Next, we examine the impact of exogenous changes in the discount factor.

Corollary 3. *(Exogenous change in δ) Suppose $\tilde{Q}$ is such that $S(\tilde{Q}) > S(Q^f)$ and $\delta \geq \underline{\delta}(\tilde{Q})$. Then, a decrease in δ has the following effects:*

1. *If $\delta \geq \underline{\delta}(\tilde{Q})$ continues to hold, then the principal continues to contract for $\tilde{Q}$ using an incomplete contract.*

2. *If $\delta < \underline{\delta}(\tilde{Q})$, then the principal contracts for a lower $\hat{Q}$ where $\delta = \underline{\delta}(\hat{Q})$ using an incomplete contract if $S(\hat{Q}) > S(Q^f)$.*

3. *If $\delta < \underline{\delta}(\tilde{Q})$, then the principal switches to a complete contract that implements Q^f if there exists no $\hat{Q}$ such that $S(\hat{Q}) > S(Q^f)$*

Proof. **Part (1)**: If $\delta \geq \underline{\delta}(\tilde{Q})$ continues to hold after an exogenous decrease in δ, then the principal continues to contract for $\tilde{Q}$ since it would remain self-enforcing.

Part (2): If $\delta < \underline{\delta}(\tilde{Q})$, then $\tilde{Q}$ is no longer self-enforcing and cannot be sustained using a relational contract. However, given the assumptions $r'(Q) > 0, r''(Q) \leq 0, c'(Q) > 0$, and $c''(Q) > 0$, we see from A.9 that $\underline{\delta}(Q)$ can be lowered by lowering Q. Therefore, for an exogenous decrease in δ, the principal has to lower her preferred quality level from $\tilde{Q}$ to some $\hat{Q}$ such that $\delta = \underline{\delta}(\hat{Q})$. $\hat{Q}$ is self-enforcing and a relational contract that implements $\hat{Q}$ will be preferred to the best complete contract that implements Q^f if $S(\hat{Q}) > S(Q^f)$.

Part (3): The proof follows the same steps as the proof for Enumerate 2 except if $S(\hat{Q}) \leq S(Q^f)$, then the principal prefers the complete contract that implements Q^f over the relational contract that implements $\hat{Q}$. □

Corollary 3 leads to Empirical Implication 3 in the main paper.

B E-Treatment Instructions

Instructions (0.80 E)

You can earn money during this experiment, with the exact amount depending on the decisions you make during the experiment. Your experimental income is calculated in points, which will be converted into cash at the rate of: $1 = 30 points. We will start you off with a balance of 150 points ($5).

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 may force us to stop the experiment.

General Information

This experiment is about how people buy and sell goods for which quality matters. Participants are divided into two groups: half will be buyers and the other half sellers. And then a trading period will start in which a buyer and seller will trade one unit of a good that can vary in quality. The price agreed upon between the buyer and seller and the quality of the good traded will determine how much money each party makes in that period. There will many trading periods throughout the course of this experiment.

Who will you trade with? At the beginning of the experiment, the computer will randomly match each participant in the room with another participant to form a buyer-seller pairing. You will be informed whether you are the buyer or seller in your pairing. You will trade with your pair-member. You will *not* be informed of the actual identity of the other person (and s/he will not be informed of your identity). *All sellers and buyers are assigned a numeric ID which is not associated with their real identity.* You will also retain your ID and role (e.g. buyer or seller) through the entire experiment.

For how many periods will you trade with the same person? All participants will remain matched with their pair-member for a random number of periods. **How is this determined?** At the end of each period, the computer will determine randomly whether the same pairings will continue for the next period or whether new pairings will be formed. In any given period, there is an **80% chance** that the same pairings will continue for the next period. In other words, in any given period, there is a 80% chance that you will continue to trade with the same person in the next period. To help you understand this, imagine that the computer has been programmed to spin a roulette wheel. If it lands on 1,2,3,4, 5, 6, 7, or 8 then you will continue to trade with the same person the next period. But if it lands on 9 or 10 the current pairings are immediately terminated. And then for the next period, the computer will randomly match you with a different person in the room to form a new pairing. This process will repeat for every new pairing. *At the beginning of each period, you will be notified on-screen whether the random matching process has kept you with the same person or matched you with a new person.*

When does the entire experiment end? If one of two conditions hold: (1) The experiment will end if all participants have already been matched with all possible trading partners. *This is because no participant will be matched with the same person more than once during this experiment.* For example, if there are 10 buyers and 10 sellers, then no buyer or seller will have more than 10 unique pairings. After 10 unique pairings, the experiment ends. (2) Even if all unique pairings have not been exhausted, the last pairing will occur once the experiment has lasted at least 18 periods. In other words, if you have traded at least 18 periods for the experiment, then your current pairing is your last one. **This does not mean the experiment stops at 18 rounds exactly;** it only means that when your last pairing **randomly** ends, you will not be paired with a new partner.

To summarize, if you have had less than 10 different trading partners during the experiment, but the experiment has not lasted at least 18 total periods, then when your current match is randomly terminated, the computer will match you with a new person and the experiment would continue. However, if the experiment has lasted at least 18 total periods, then the experiment will end once your current pairing is randomly terminated.

CONDUCTING TRADES

Each trade occurs within a trading period. Each trading period is then divided into a *proposal phase* followed by *a quality determination phase* and then followed by a *payment determination phase.*

a) During the *proposal phase*, the buyer can make a proposal on the terms of trade to the seller. The seller can either accept or reject the proposal.
b) If the seller accepts the proposal, then during the *quality determination phase*, the seller chooses the actual quality level to supply.
c) After quality is observed, comes the *payment determination phase*. During this phase, the buyer can make final adjustments in payment depending on the initial terms of the proposal.

During each phase, you can take as much time as you need to make a good decision, but the faster you make your decision, the faster the experiment will move.

Specific details of each phase are given below:

1. The Proposal Phase

Each period starts with a proposal phase. A proposal allows the parties to agree to the terms of trade by including a list of promises and obligations of both parties (see below for details). *The buyer can submit a single proposal during the proposal phase. Once a proposal is submitted, the seller will decide to accept or reject the proposal.*

How does a buyer make a proposal? A proposal screen will appear that will require the buyer to enter values for the following terms: *desired quality*, *price*, and a *performance bonus*. These terms are described below.

a) **Desired quality** – The buyer must (1) ask the seller to deliver a specific quality level and (2) specify whether the quality level is binding or discretionary (if binding, the computer enforces the quality level).

   Regarding (1), possible quality levels can range from 1 to 15, where higher numbers indicate higher quality (whole numbers only). Buyers earn more when they get higher quality. The buyer should enter a number between 1 to 15 in the "Desired quality" field.

   Regarding (2), The buyer also specifies whether s/he wants desired quality to be **binding** or **discretionary** by clicking the appropriate checkbox. **Binding** is similar to a legally binding obligation – once the seller agrees to the proposal, the computer will ensure that the seller supplies the desired quality level. **Discretionary** means that the obligation is informal rather than legal – i.e. the seller's quality choice will not be enforced by the computer. Thus, nothing restricts the seller from choosing a quality level that is different from the desired quality during the quality determination phase.

b) **Price** – This allows the buyer to state the price she will pay for the good. The buyer enters a price in the "*Price*" field. The price ranges from 0 to 200 (whole numbers).

   The price the buyer specifies will be *binding*. It is similar to a legally binding obligation – once the proposal is agreed upon, the computer will ensure that the price is paid to the seller.

c) **Performance bonus**– For the case when desired quality is discretionary, the buyer can state that s/he will pay a bonus that might be linked to quality. *To enter a bonus, click on the "yes"* *box next to "would you like to offer a bonus?"* Then enter a number in in the "*Bonus*" field to specify the size of the bonus (enter a whole number from 0 to 200). *If the buyer does not wish to offer a bonus, simply click "no" next to "would you like to offer a bonus?"* The total payment is *price* plus *bonus*.

Important: The stated bonus is *not binding*. During the payment determination phase to come later, the buyer can choose any bonus level s/he wishes. Thus, this is a discretionary bonus. However, if the buyer clicked "no" to offering bonus, then there will be no payment determination phase for the buyer in this period. The **Price** then becomes the final payment.

After the buyer has specified desired quality, whether quality is binding or not, price, and performance Bonus (if quality is discretionary), s/he needs to click "OK" to submit the offer. Next comes the quality determination phase.

2. Quality Determination Phase

Following the proposal phase, all sellers who accepted an agreement that did not have a binding *Desired quality* will determine the level of quality that they will supply to their buyers. A seller can choose any quality s/he wants to from 1 to 15. The *Quality Determination* Screen will appear and a seller can enter his/her quality choice in the "*Actual Quality"* field. Nothing restricts the seller from choosing a quality level that is different from the "desired quality" level specified in the proposal. **Note:** If the buyer chose a binding quality, then there is no quality determination phase for the seller.

3. Payment Determination Phase

Following the quality determination phase, all buyers who offered a bonus will determine the level of actual bonus that s/he will pay to the seller. During this phase, **after discretionary actual quality is observed by the buyer**, the buyer will choose actual bonus to be paid to the seller. The *Payment Determination* screen will appear and the buyer will enter his/her bonus choice in the "*Actual Bonus"* field. **Nothing restricts the buyer from choosing a bonus level that is different from the bonus that was specified in the proposal**. The actual bonus can range from 0 to 200 at the buyer's discretion.

Note: If the buyer chose a binding quality or did not offer a bonus, then there is no payment determination phase for the buyer.

How Are Points (Income) Calculated?

How do Buyers Make Money?

- If the **buyer does not make an offer or the seller rejects the offer**, the **buyer will receive 15** points for that period.
- If the buyer's proposal is accepted, the buyer's points for the period depend on the actual quality, the price and the actual bonus paid. That is,

Buyer Points = 12*Actual Quality – Price –Actual Bonus

- As you can see, the higher the actual quality, the more points the buyer earns. At the same time, the lower total payments (price plus actual bonus), the more points the buyer earns.
- **In summary, higher quality at lower payments means more points for the buyer.**

How do Sellers Make Money?

- If the **seller rejects the proposal or the buyer does not make an offer**, the **seller will receive 15** points for that period.

- If the seller has accepted an offer, then the seller's points depends on the price, actual bonus, and production costs s/he incurs. The points of a seller is determined as follows:

Seller Points = Price +Actual Bonus– Production Costs

- As you can see, **the higher the actual payments, 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. Roughly speaking, the cost is determined by the following formula: $Cost = \frac{q^2}{2}$. We say "roughly speaking" because we will round the cost number to the nearest whole number. The following table gives you the exact cost in whole numbers of producing each quality level.

Quality	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Cost	**1**	**2**	**5**	**8**	**13**	**18**	**25**	**32**	**41**	**50**	**61**	**72**	**85**	**98**	**113**

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.** Note that buyers and sellers can incur losses in each period. These losses are subtracted from your points balance.

At the end of each period, the buyer and seller will be shown an "**income screen**." The following information is displayed on this screen:

- the ID number of your trading partner.
- the **Price** the buyer offered.
- the **Proposed Bonus**
- the **Actual bonus** granted
- the buyer's **Desired Quality and whether it was binding or not.**
- the **Actual quality** delivered by the seller.
- the points earned (lost) by both parties 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 so that you can learn from your past results.

At the beginning of the next period, the computer will inform you if you have been randomly matched with the same trading partner or with a different partner.

Before we begin the experiment, we ask all participants to complete a questionnaire which will test 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 proposal phase** so that you can get accustomed to the computer. During the trial periods, no money can be earned. Your ID numbers will also be suppressed on the screen during the trial periods.

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, a desired bonus of 20, and indicated a non-binding desired quality of 9. A seller accepts your offer and actually chooses a quality of 8. You pay an actual bonus of 10. How many points did you earn for this period?
3. Suppose that you are a buyer and you offered a price of 70, a desired bonus of 10, and indicated a non-binding desired quality of 10. A seller accepts your offer and chooses actual quality of 10. If you choose to pay an actual bonus of 10, how many points did you earn for this period?
4. Suppose that you are a seller and you just finished trading with buyer no. 3. What is the probability that you will not trade with buyer no. 3 the next period?
5. Suppose that you are a seller and you did not accept an offer during the trading phase. How many points do you earn for this period?
6. (True or false) Suppose that you are a buyer and you have already finished 19 trading periods for the experiment across four different sellers. Once your relationship with your current trading partner is terminated, will you be paired with another seller.
7. Suppose that you are a seller and that you accepted an offer with a price of 40, a non-binding desired quality of 2, and a desired bonus of 5. You choose to supply an actual quality of 5. If your buyer pays you an actual bonus of 10, how many points did you earn for this period?
8. Suppose that you are a buyer and you offered a proposal with a binding desired quality of 5. The actual quality chosen by the seller must be what?
9. Suppose that you are a seller and you accepted a proposal with a binding Desired quality of 10 and a Price of 50. How many points did you earn this period?

Answers

1. 15
2. 56
3. 40
4. 20%
5. 15
6. False. The experiment will end.
7. 37
8. 5 The seller cannot deviate from 5 when quality is binding.
9. 0.

C PE0.80-Treatment Instructions

Instructions (0.80 PE)

You can earn money during this experiment, with the exact amount depending on the decisions you make during the experiment. Your experimental income is calculated in points, which will be converted into cash at the rate of: $1 = 30 points. We will start you off with a balance of 150 points ($5).

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 may force us to stop the experiment.

General Information

This experiment is about how people buy and sell goods for which quality matters. Participants are divided into two groups: half will be buyers and the other half sellers. And then a trading period will start in which a buyer and seller will trade one unit of a good that can vary in quality. The price agreed upon between the buyer and seller and the quality of the good traded will determine how much money each party makes in that period. There will many trading periods throughout the course of this experiment.

Who will you trade with? At the beginning of the experiment, the computer will randomly match each participant in the room with another participant to form a buyer-seller pairing. You will be informed whether you are the buyer or seller in your pairing. You will trade with your pair-member. You will *not* be informed of the actual identity of the other person (and s/he will not be informed of your identity). *All sellers and buyers are assigned a numeric ID which is not associated with their real identity.* You will also retain your ID and role (e.g. buyer or seller) through the entire experiment.

For how many periods will you trade with the same person? All participants will remain matched with their pair-member for a random number of periods. **How is this determined?** At the end of each period, the computer will determine randomly whether the same pairings will continue for the next period or whether new pairings will be formed. In any given period, there is an **80% chance** that the same pairings will continue for the next period. In other words, in any given period, there is a 80% chance that you will continue to trade with the same person in the next period. To help you understand this, imagine that the computer has been programmed to spin a roulette wheel. If it lands on 1,2,3,4, 5, 6, 7, or 8 then you will continue to trade with the same person the next period. But if it lands on 9 or 10 the current pairings are immediately terminated. And then for the next period, the computer will randomly match you with a different person in the room to form a new pairing. This process will repeat for every new pairing. *At the beginning of each period, you will be notified on-screen whether the random matching process has kept you with the same person or matched you with a new person.*

When does the entire experiment end? If one of two conditions hold: (1) The experiment will end if all participants have already been matched with all possible trading partners. *This is because no participant will be matched with the same person more than once during this experiment.* For example, if there are 10 buyers and 10 sellers, then no buyer or seller will have more than 10 unique pairings. After 10 unique pairings, the experiment ends. (2) Even if all unique pairings have not been exhausted, the last pairing will occur once the experiment has lasted at least 18 periods. In other words, if you have traded at least 18 periods for the experiment, then your current pairing is your last one. **This does not mean the experiment stops at 18 rounds exactly;** it only means that when your last pairing **randomly** ends, you will not be paired with a new partner.

To summarize, if you have had less than 10 different trading partners during the experiment, but the experiment has not lasted at least 18 total periods, then when your current match is randomly terminated, the computer will match you with a new person and the experiment would continue. However, if the experiment has lasted at least 18 total periods, then the experiment will end once your current pairing is randomly terminated.

CONDUCTING TRADES

Each trade occurs within a trading period. Each trading period is then divided into a *proposal phase* followed by *a quality determination phase* and then followed by a *payment determination phase.*

a) During the *proposal phase*, the buyer can make a proposal on the terms of trade to the seller. The seller can either accept or reject the proposal.
b) If the seller accepts the proposal, then during the *quality determination phase*, the seller chooses the actual quality level to supply.
c) After quality is observed, comes the *payment determination phase*. During this phase, the buyer can make final adjustments in payment depending on the initial terms of the proposal.

During each phase, you can take as much time as you need to make a good decision, but the faster you make your decision, the faster the experiment will move.

Specific details of each phase are given below:

1. The Proposal Phase

Each period starts with a proposal phase. A proposal allows the parties to agree to the terms of trade by including a list of promises and obligations of both parties (see below for details). *The buyer can submit a single proposal during the proposal phase. Once a proposal is submitted, the seller will decide to accept or reject the proposal.*

How does a buyer make a proposal? A proposal screen will appear that will require the buyer to enter values for the following terms: *desired quality*, *price*, and a *performance bonus*. These terms are described below.

a) **Desired quality** – The buyer must (1) ask the seller to deliver a specific quality level and (2) specify whether the quality level is binding or discretionary (if binding, the computer enforces the quality level).

   Regarding (1), possible quality levels can range from 1 to 15, where higher numbers indicate higher quality (whole numbers only). Buyers earn more when they get higher quality.

   Regarding (2), The buyer also specifies whether s/he wants desired quality to be **binding** or **discretionary** by clicking the appropriate checkbox. **Binding** is similar to a legally binding obligation – once the seller agrees to the proposal, the computer will ensure that the seller supplies the desired quality level. **Discretionary** means that the obligation is informal rather than legal – i.e. the seller's quality choice will not be enforced by the computer. Thus, nothing restricts the seller from choosing a quality level that is different from the desired quality during the quality determination phase. **However, not all quality levels can be made binding. Only quality levels "1" and "5" can be made binding.**

   Therefore, if the buyer clicks "*binding*", then s/he must also click "1" or "5" in *Desired quality* checkbox right next to the "binding" checkbox.

   If the buyer clicks "*discretionary*", then s/he must enter a number between 1 to 15 in the field next to the *discretionary* checkbox.

b) **Price** – This allows the buyer to state the price she will pay for the good. The buyer enters a price in the "*Price*" field. The price ranges from 0 to 200 (whole numbers).

The price the buyer specifies will be *binding*. It is similar to a legally binding obligation – once the proposal is agreed upon, the computer will ensure that the price is paid to the seller.

c) **Performance bonus**– For the case when desired quality is discretionary, the buyer can state that s/he will pay a bonus that might be linked to quality. *To enter a bonus, click on the "yes" box next to "would you like to offer a bonus?"* Then enter a number in in the "*Bonus*" field to specify the size of the bonus (enter a whole number from 0 to 200). *If the buyer does not wish to offer a bonus, simply click "no" next to "would you like to offer a bonus?"* The total payment is *price* plus *bonus*.

Important: The stated bonus is *not binding*. During the payment determination phase to come later, the buyer can choose any bonus level s/he wishes. Thus, this is a discretionary bonus. However, if the buyer clicked "no" to offering bonus, then there will be no payment determination phase for the buyer in this period. The **Price** then becomes the final payment.

After the buyer has specified desired quality, price and performance bonus, s/he needs to click "OK" to submit it. Next comes the quality determination phase.

2. Quality Determination Phase

Following the proposal phase, all sellers who accepted an agreement that did not have a binding *Desired quality* level of "1" or "5" will determine the level of quality that they will supply to their buyers. A seller can choose any quality s/he wants to from 1 to 15. The *Quality Determination* Screen will appear and a seller can enter his/her quality choice in the "*Actual Quality"* field. Nothing restricts the seller from choosing a quality level that is different from the "desired quality" level specified in the proposal.

Note: If the buyer chose a binding quality of "1" or "5", then there is no quality determination phase for the seller.

3. Payment Determination Phase

Following the quality determination phase, all buyers who offered a bonus will determine the level of actual bonus that s/he will pay to the seller. During this phase, **after quality is observed by the buyer**, the buyer will choose actual bonus to be paid to the seller. The *Payment Determination* screen will appear and the buyer will enter his/her bonus choice in the "*Actual Bonus"* field. **Nothing restricts the buyer from choosing a bonus level that is different from the bonus that was specified in the proposal**. The actual bonus can range from 0 to 200 at the buyer's discretion.

How Are Points (Income) Calculated?

How do Buyers Make Money?

- If the **buyer does not make an offer or the seller rejects the offer**, the **buyer will receive 15** points for that period.
- If the buyer's proposal is accepted, the buyer's points for the period depend on the actual quality, the price and the actual bonus paid. That is,

Buyer Points = 12*Actual Quality – Price –Actual Bonus

- As you can see, the higher the actual quality, the more points the buyer earns. At the same time, the lower total payments (price plus actual bonus), the more points the buyer earns.
- **In summary, higher quality at lower payments means more points for the buyer.**

How do Sellers Make Money?

- If the **seller rejects the proposal or the buyer does not make an offer**, the **seller will receive 15** points for that period.
- If the seller has accepted an offer, then the seller's points depends on the price, actual bonus, and production costs s/he incurs. The points of a seller is determined as follows:

Seller Points = Price +Actual Bonus– Production Costs

- As you can see, **the higher the actual payments, 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. Roughly speaking, the cost is determined by the following formula: $Cost = \frac{q^2}{2}$. We say "roughly speaking" because we will round the cost number to the nearest whole number. The following table gives you the exact cost in whole numbers of producing each quality level.

Quality	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Cost	**1**	**2**	**5**	**8**	**13**	**18**	**25**	**32**	**41**	**50**	**61**	**72**	**85**	**98**	**113**

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.** Note that buyers and sellers can incur losses in each period. These losses are subtracted from your points balance.

At the end of each period, the buyer and seller will be shown an "**income screen**." The following information is displayed on this screen:

- the ID number of your trading partner.
- the **Price** the buyer offered.
- the **Proposed Bonus**
- the **Actual bonus** granted
- the buyer's **Desired Quality and whether it was binding or not.**
- the **Actual quality** delivered by the seller.
- the points earned (lost) by both parties 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 so that you can learn from your past results.

At the beginning of the next period, the computer will inform you if you have been randomly matched with the same trading partner or with a different partner.

Before we begin the experiment, we ask all participants to complete a questionnaire which will test 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 proposal phase** so that you can get accustomed to the computer. During the trial periods, no money can be earned. Your ID numbers will also be suppressed on the screen during the trial periods.

Instructions (0.80 PE)

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, a desired bonus of 20, and indicated a desired quality of 9. A seller accepts your offer and actually chooses a quality of 8. You pay an actual bonus of 10. How many points did you earn for this period?
3. Suppose that you are a buyer and you offered a price of 70, a desired bonus of 10, and indicated a desired quality of 10. A seller accepts your offer and chooses actual quality of 10. If you choose to pay an actual bonus of 10, how many points did you earn for this period?
4. Suppose that you are a seller and you just finished trading with buyer no. 3. What is the probability that you will not trade with buyer no. 3 the next period?
5. Suppose that you are a seller and you did not accept an offer during the trading phase. How many points do you earn for this period?
6. (True or false) Suppose that you are a buyer and you have already finished 19 trading periods for the experiment across four different sellers. Once your relationship with your current trading partner is terminated, will you be paired with another seller.
7. Suppose that you are a seller and that you accepted an offer with a price of 40, a desired quality of 2, and a desired bonus of 5. You choose to supply an actual quality of 5. If your buyer pays you an actual bonus of 10, how many points did you earn for this period?
8. Suppose that you are a buyer and you offered a proposal with a binding desired quality of 5. The actual quality chosen by the seller must be what?
9. Suppose that you are a seller and you accepted a proposal with a Desired quality of 4. Can you deviate from 4 in the quality determination phase?

Answers

1. 15
2. 56
3. 40
4. 20%
5. 15
6. False. The experiment will end.
7. 37
8. 5 The seller cannot deviate from 5 when 5 is binding. Remember that the buyer can make quality levels of 1 or 5 binding.
9. Yes. The only quality levels that can be made binding are 1 and 5.

Note: PE0.50 instructions are identical to the PE0.80 instructions except the probability of trading with the same person the following period is 50% instead of 80%

D Screenshots for E Treatment

This section contains the screenshots for the E treatment.

One minor point is that all the screens show the remaining time on the upper right hand corner. In the decision screens, in which subjects had to take an action (e.g. contract formation, quality determination, bonus determination), subjects had five minutes to make a decision. This is a generous amount of time and only a few outliers ran out of time and usually near the beginning of a session.

The screen shots are presented in the same order as the sequence of moves within a stage-game.

Period starts with buyer offer screen

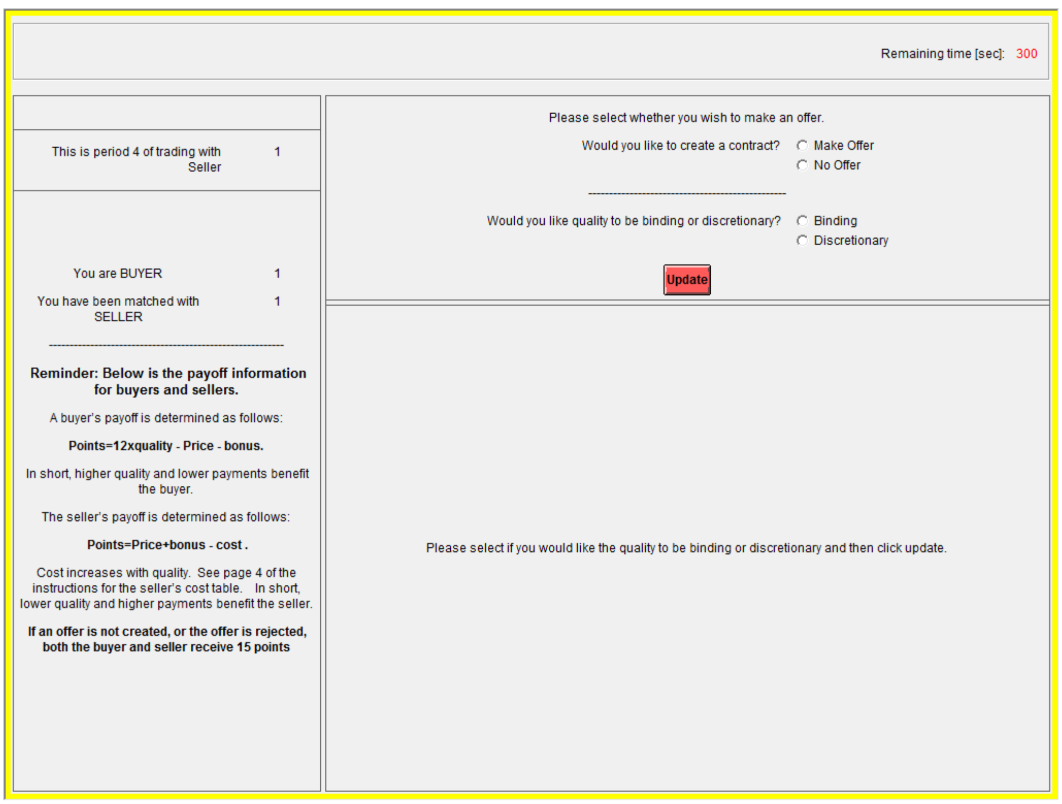

If buyer chooses "No offer," and clicks "Update", this is what s/he sees

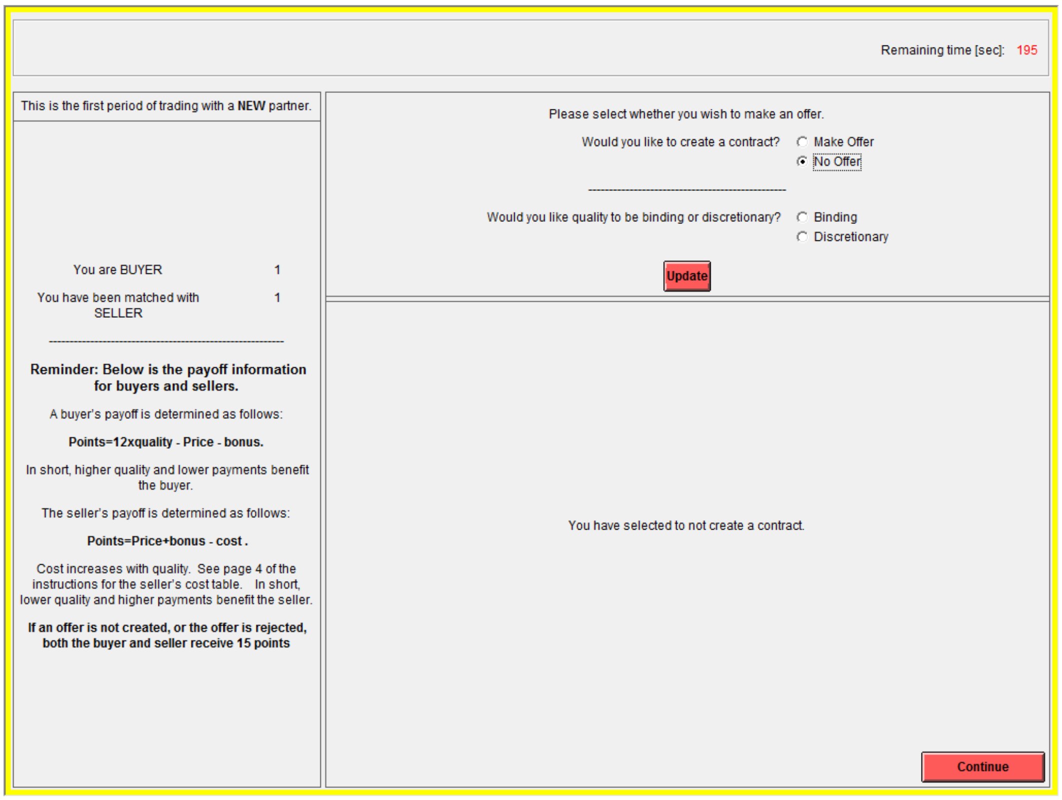

After pressing "Continue" on the previous screen, the subjects are shown the following end of period summary screen

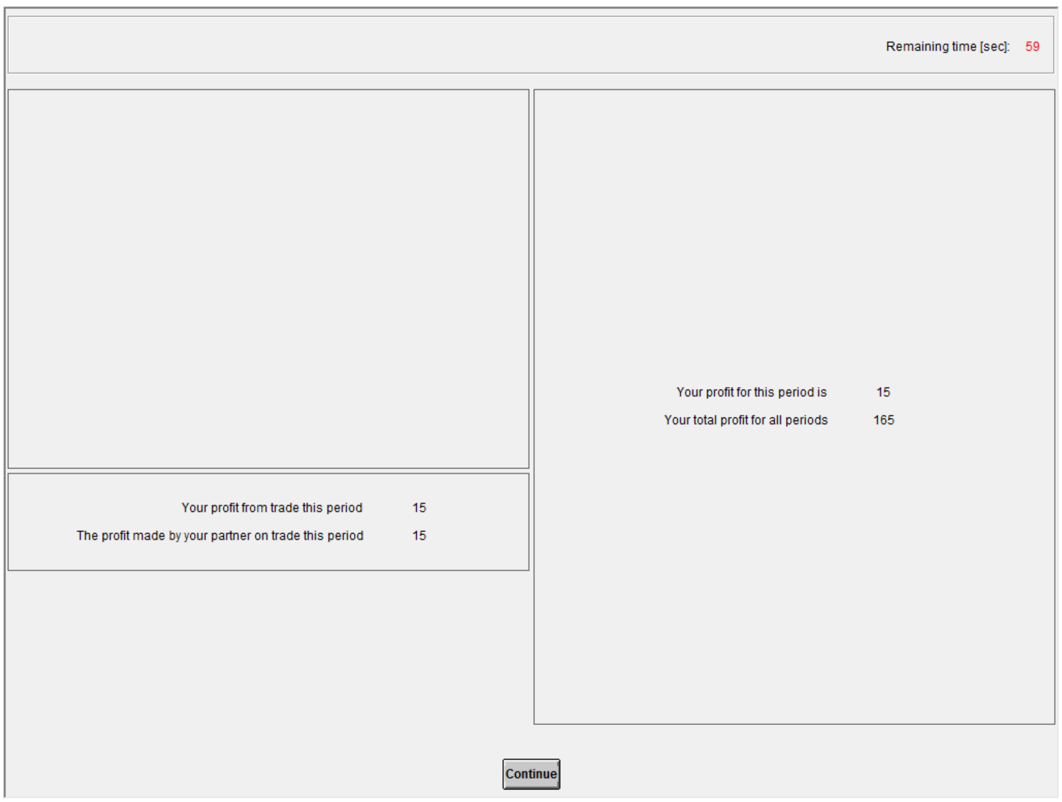

If instead the buyer had clicked “Make Offer” and “Binding” to create a binding contract that enforces quality and price, then the buyer offer screen (after clicking “Update”) changes to the screen below. Note the buyer must select the binding quality level and enter an offered Price.

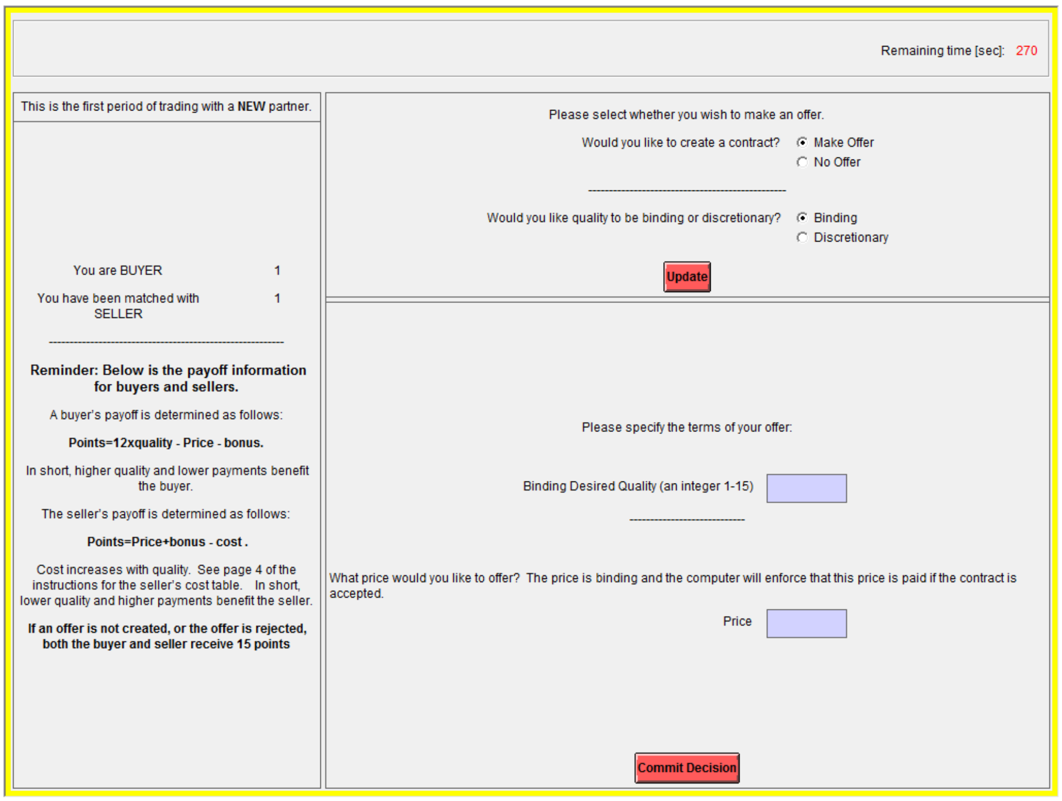

Suppose we enter binding quality of 12 and a price of 50. Then pressing "Commit Decision" takes us to the next screen for the buyer. The buyer waits at this screen because the seller must decide whether to accept or reject the contract. Note that the default bonus for a binding contract is 0 since the bonus plays no incentive role in a binding simple contract.

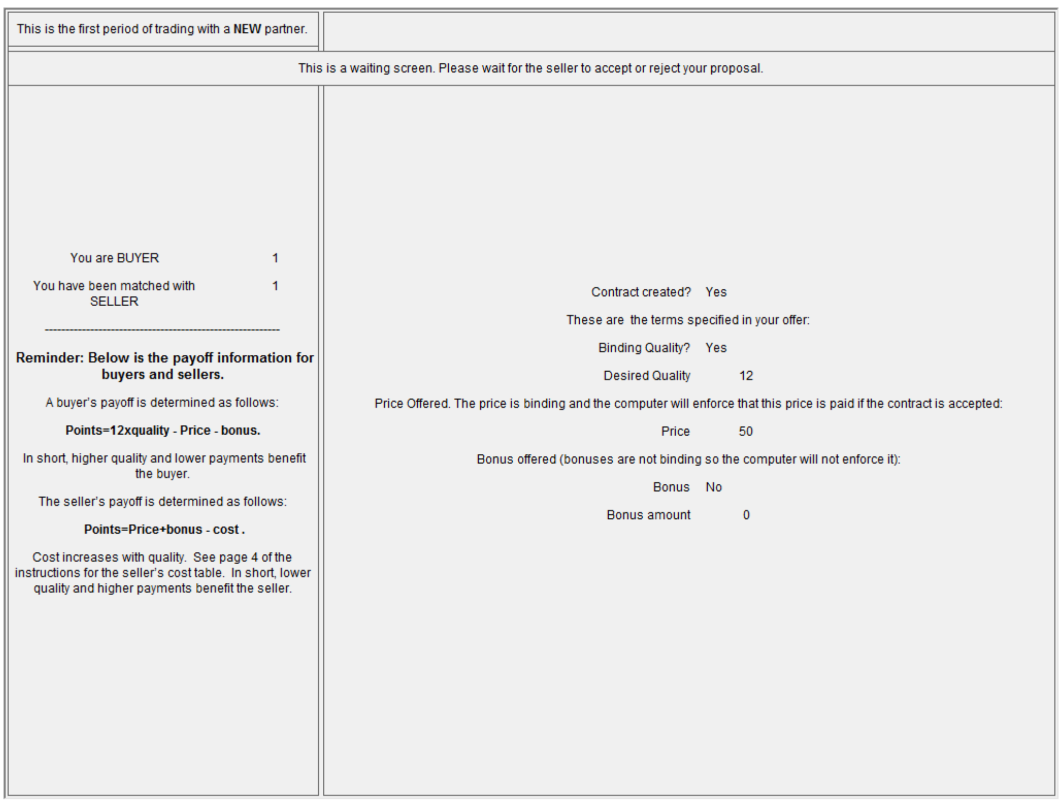

While the buyer is waiting, the seller sees the following screen.

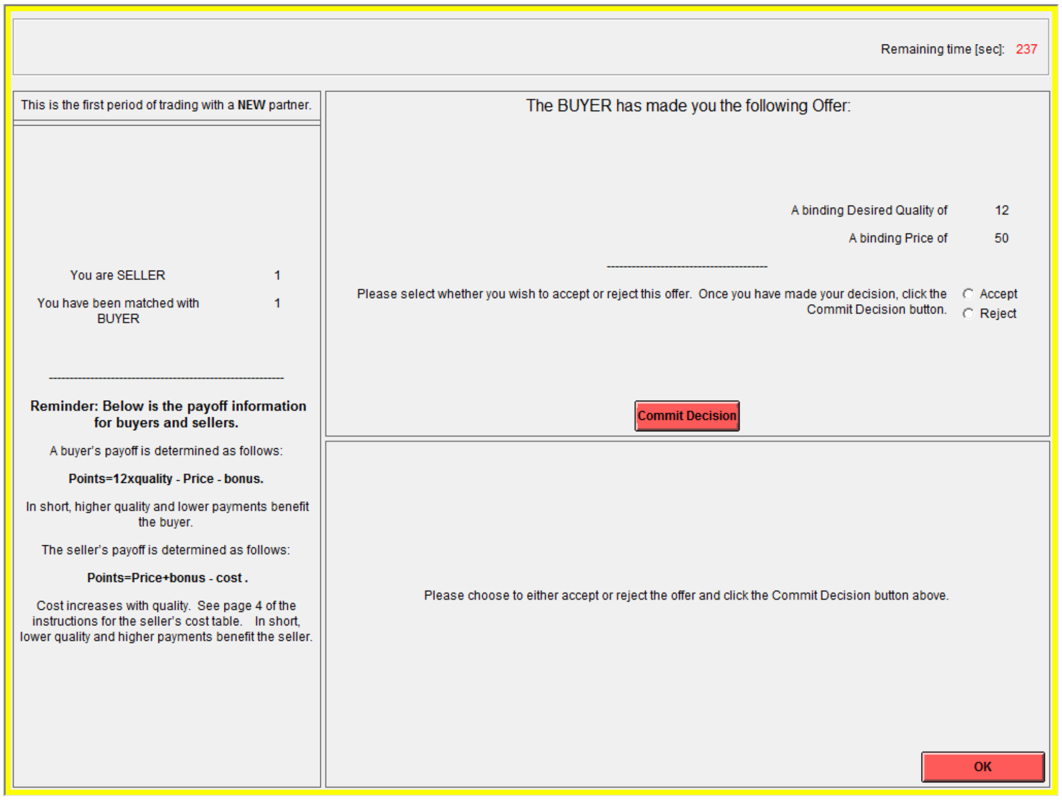

If the seller rejects the contract, then the seller is taken to the following screen (the buyer is shown an analogous screen)

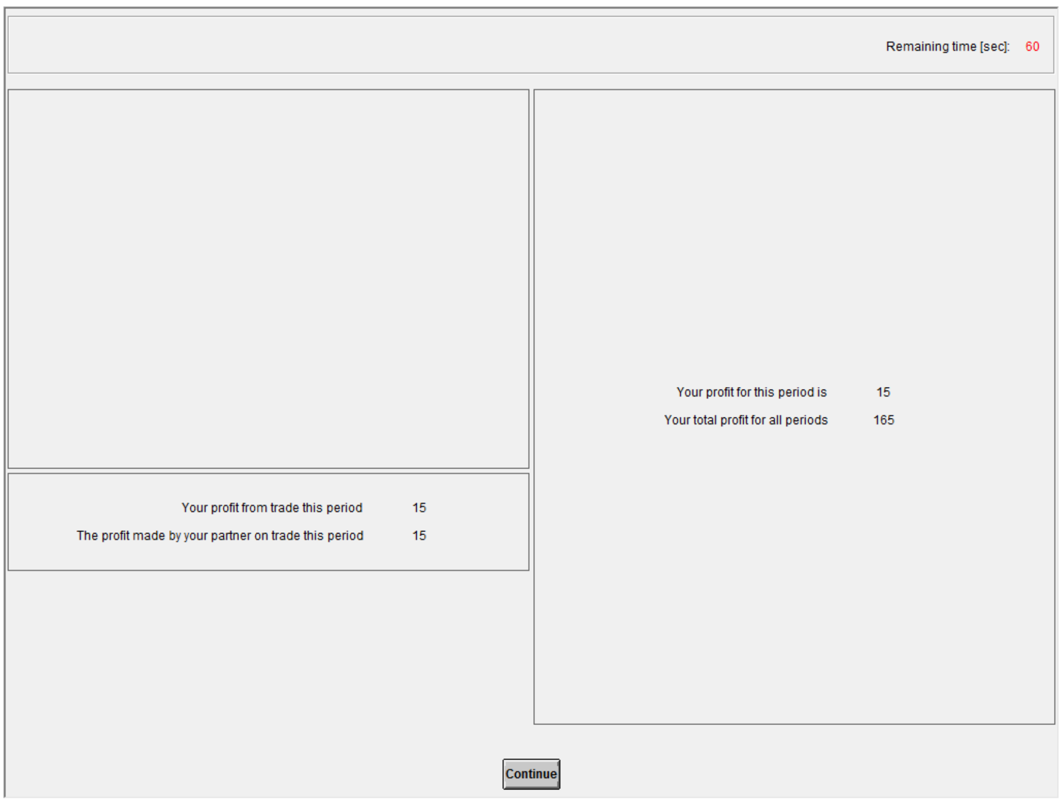

If the seller instead accepts the contract, then the trade is completed (there is no ex post discretion to choose quality or payments under a binding contract) and taken to the following screen (the buyer is shown an analogous screen)

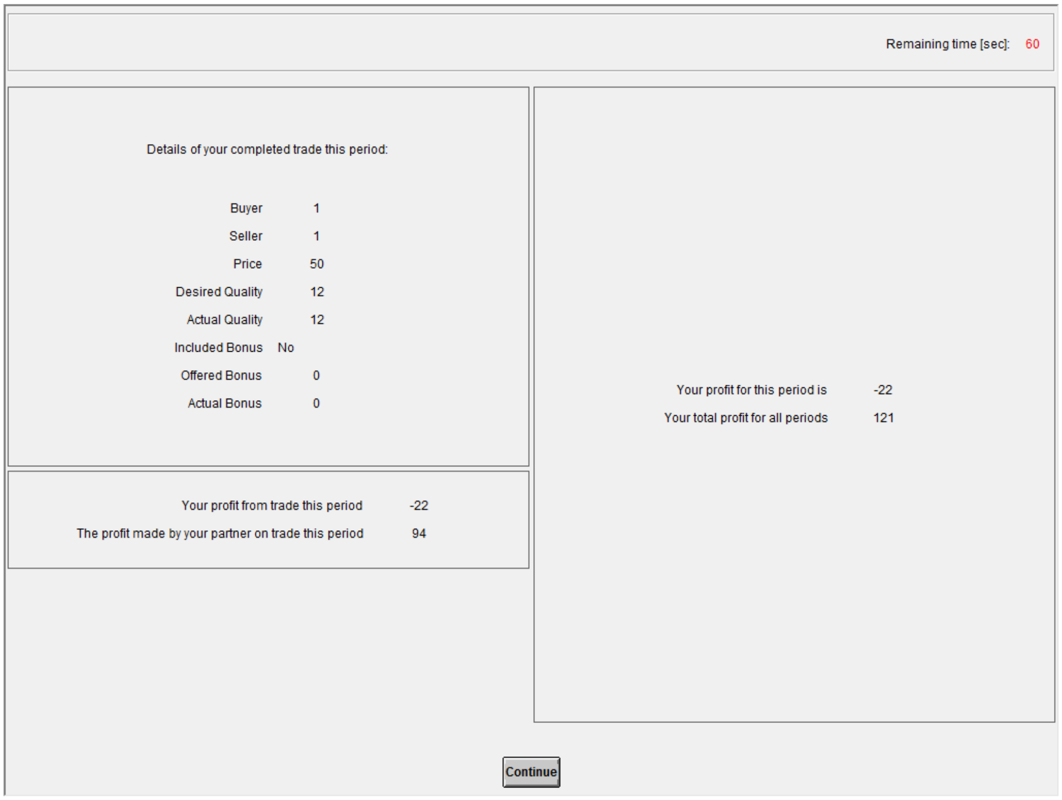

Now let's return to the buyer offer screen. Had the buyer chosen a discretionary contract, then the offer screen changes to the following:

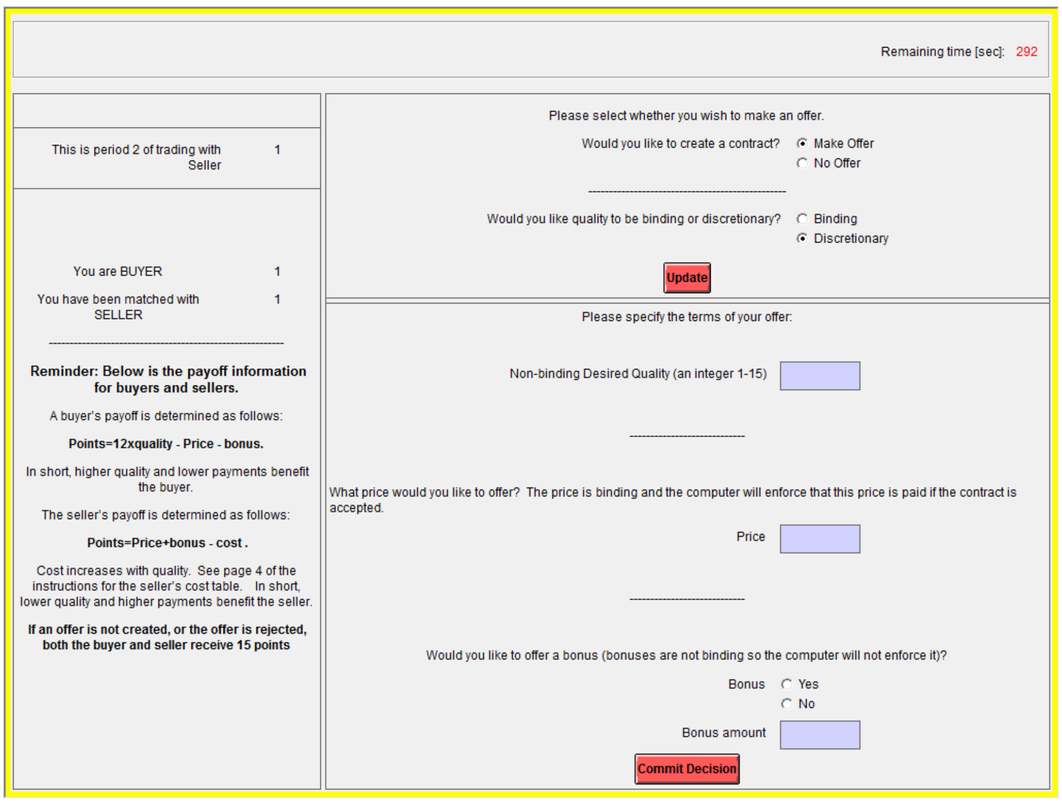

If the buyer offers a discretionary contract asking for Q=8, P=30 and B=40, then after clicking “Commit Decision” s/he is taken to the following waiting screen while the seller is making an accept or reject decision.

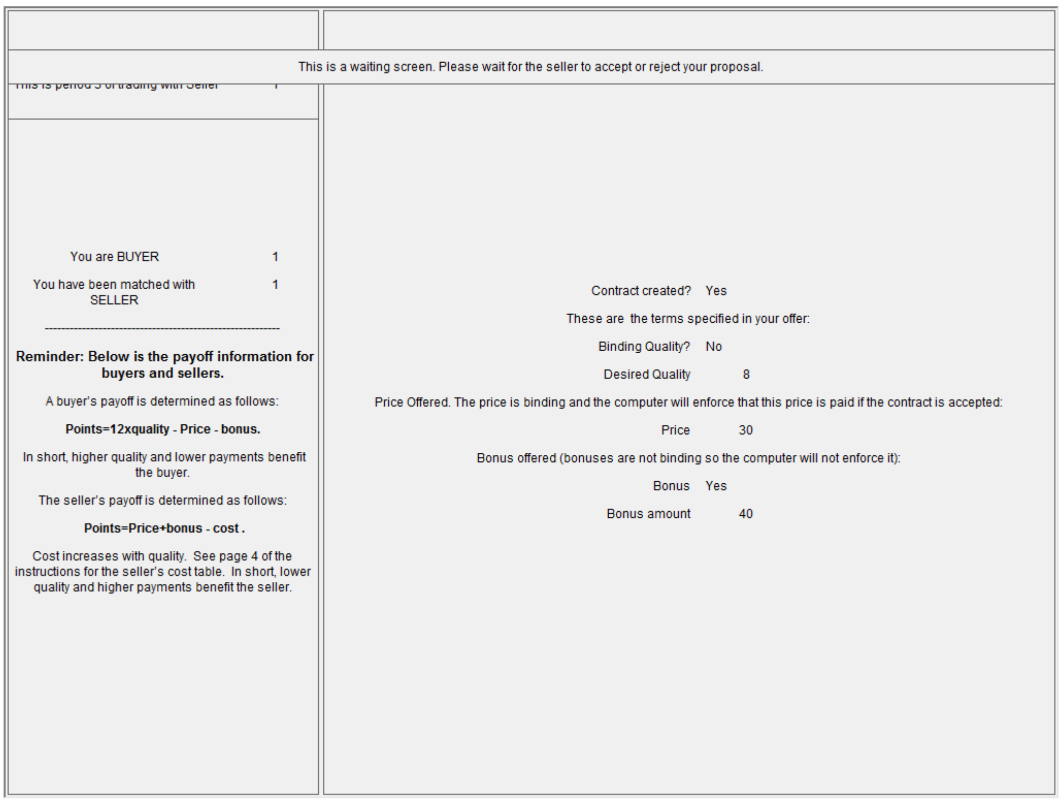

If the seller rejects the discretionary contract, then both buyer and seller are taken to the end of the period screen much like what has already been shown earlier. However, if the seller accepts the contract, her decision screen looks like the following (note once s/he chooses accept, a quality determination box appears at the bottom of the screen):

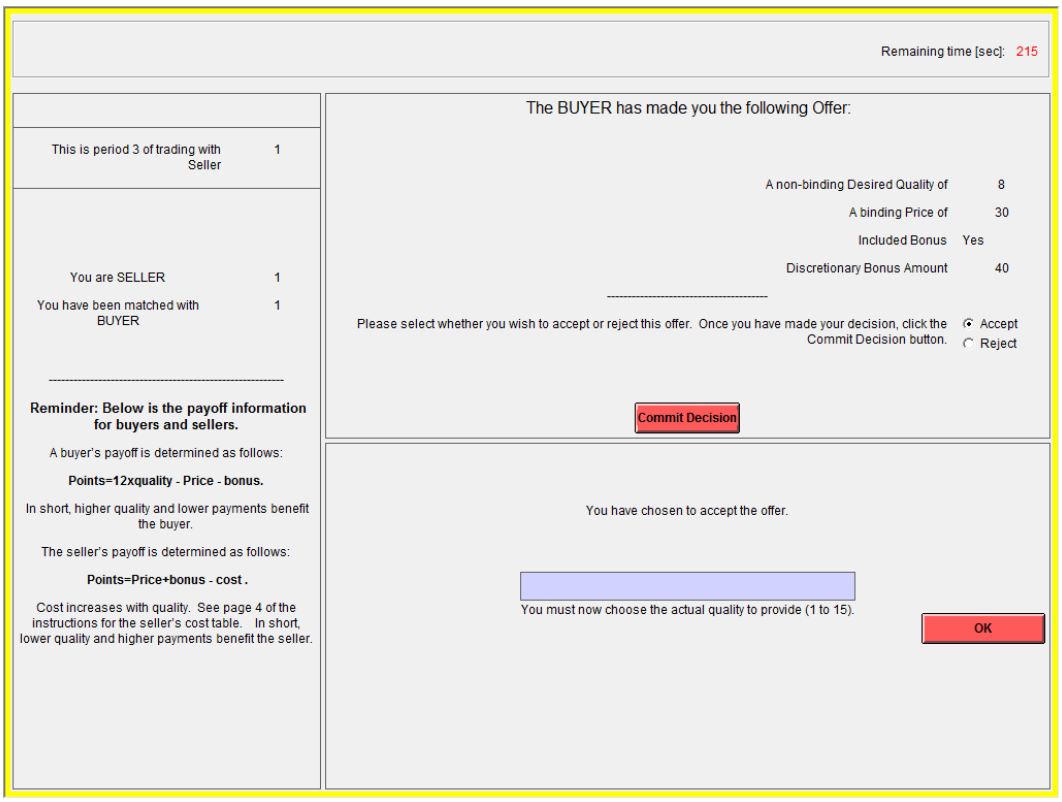

If the seller chooses an actual quality of q=9, s/he is taken to the following waiting screen.

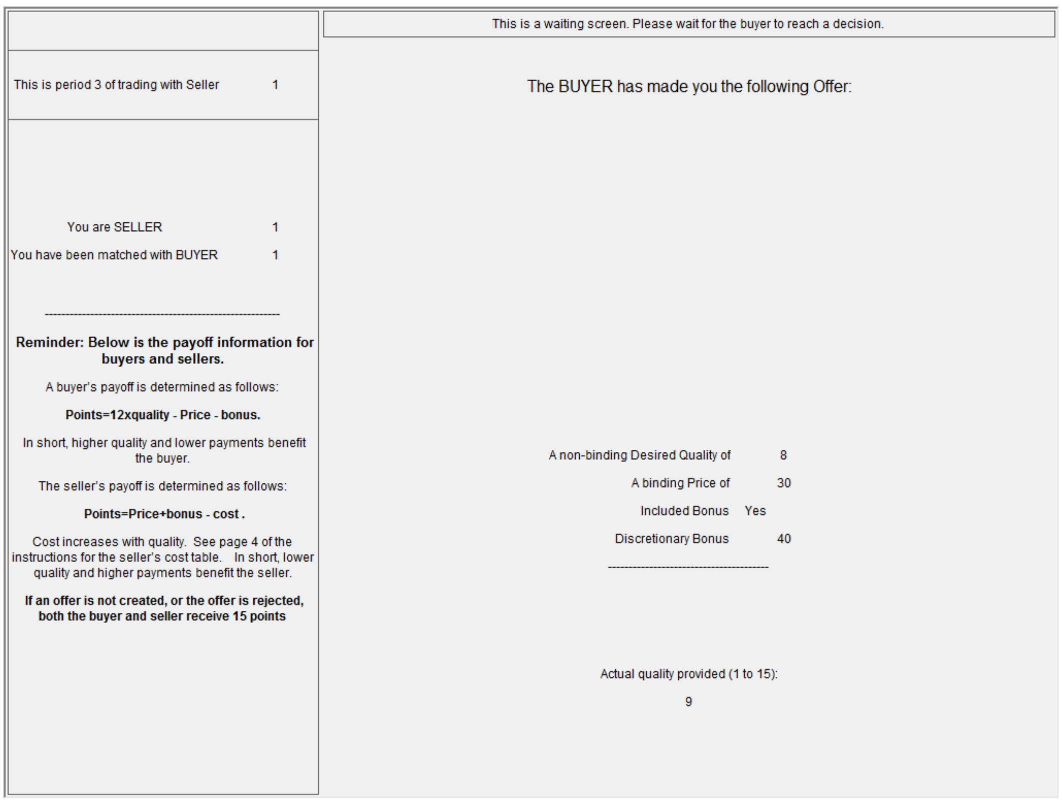

While the seller is waiting, the buyer is taken to the following bonus determination screen.

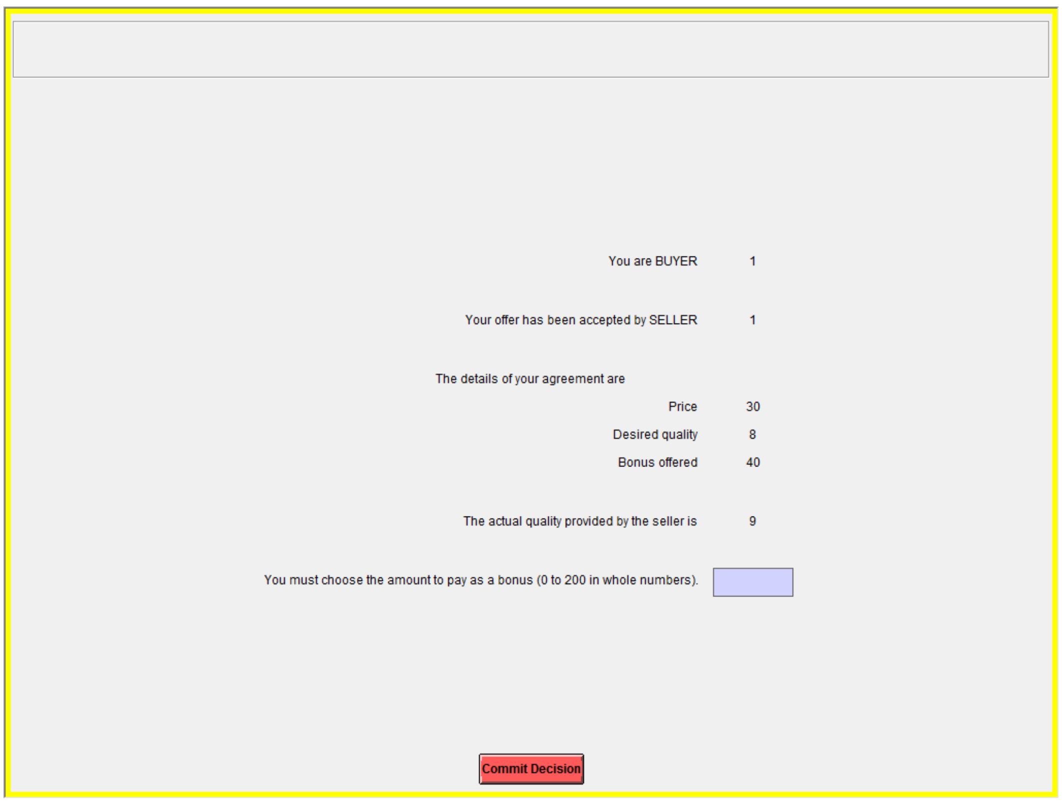

If the buyer pays an actual bonus of b=40 and then presses "Commit Decision," s/he is taken to the following end of the period summary screen. The seller sees an analogous screen.

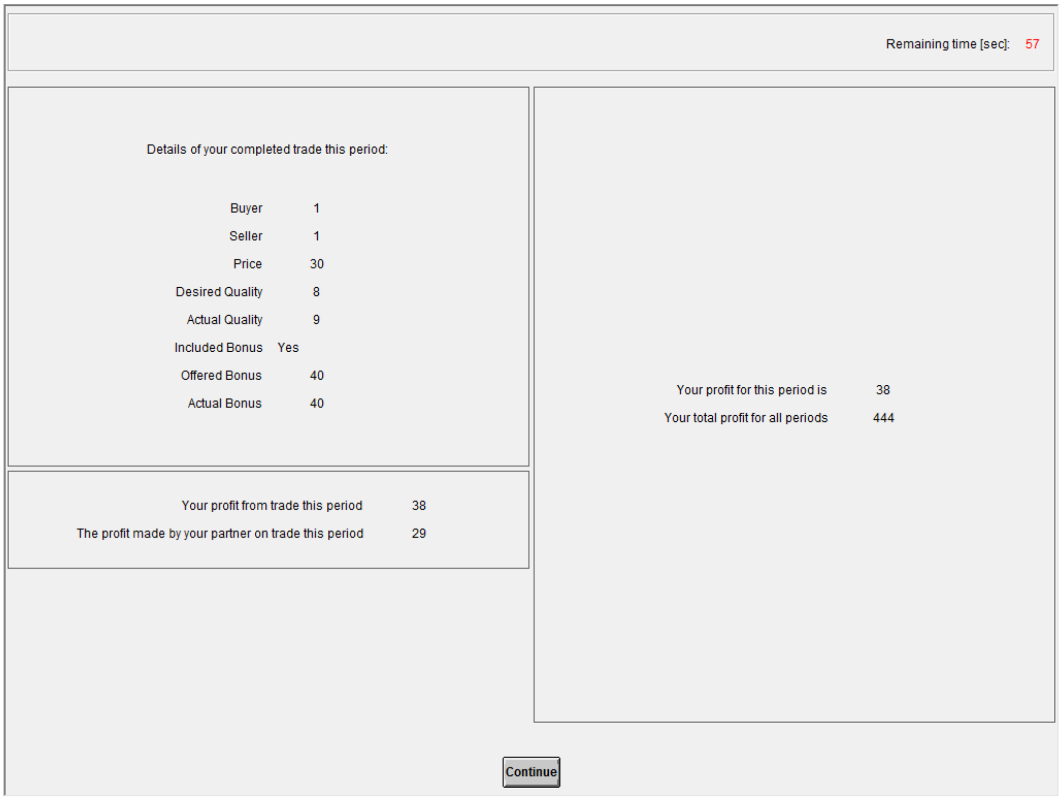

Once a period is over, both the buyer and seller see the following screen that determines their probability of trading with each other again the next period. A key point to note is that, as a practical matter, the realized draw of the continuation probability was simultaneously applied to all pairs of buyers and sellers in a session to facilitate orderly rematching when supergames terminate. In other words, either all pairs in the room continued or terminated in the same period. This made it easy to implement stranger matching. Nonetheless, to ensure saliency of the continuation probability, we forced each subject to press the "Reveal Draw" button to show them the realized draw (whether they will be rematched with the same partner or a new partner). To speed up the experiment, they were given a maximum of 15 seconds to press the button. After 15 seconds, the next period begins and the buyer offer screen appears. The experimenter announced to subjects whether they are rematched with the same person or matched with a new person. Moreover, the top left side of the decision screens for both the buyer and seller remind them how many periods they have been trading with the same partner. Thus, even if some subjects forgot to press the "Reveal Draw" button, subjects were still informed of the realized draw because we implemented multiple layers of prompts to ensure that subjects knew the draw.

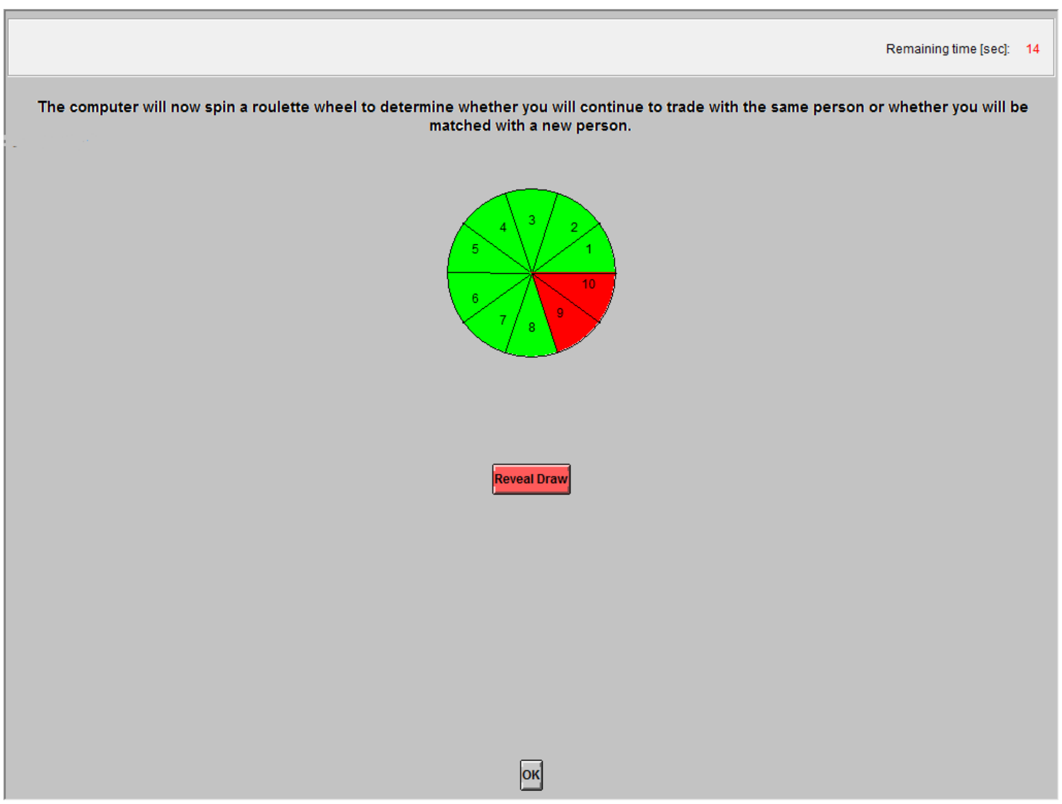

The next screen shows the revealed draw after a subject presses the "Reveal Draw" button

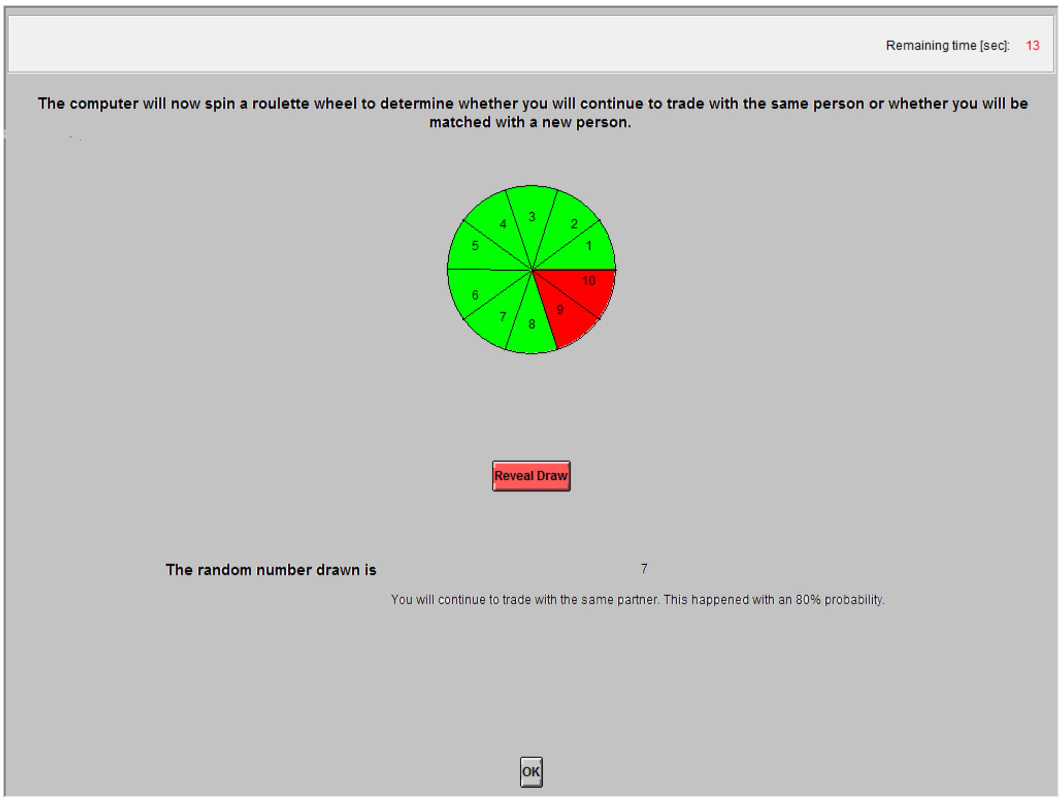

E Screenshots for PE0.80 Treatment

This section contains the screenshots for the PE0.80 treatment. PE0.50 screenshots were not included because they are identical except for the roulette wheel that determines the probability of continuation.

The screen shots are presented in the same order as the sequence of moves within a stage-game.

Period starts with buyer offer screen

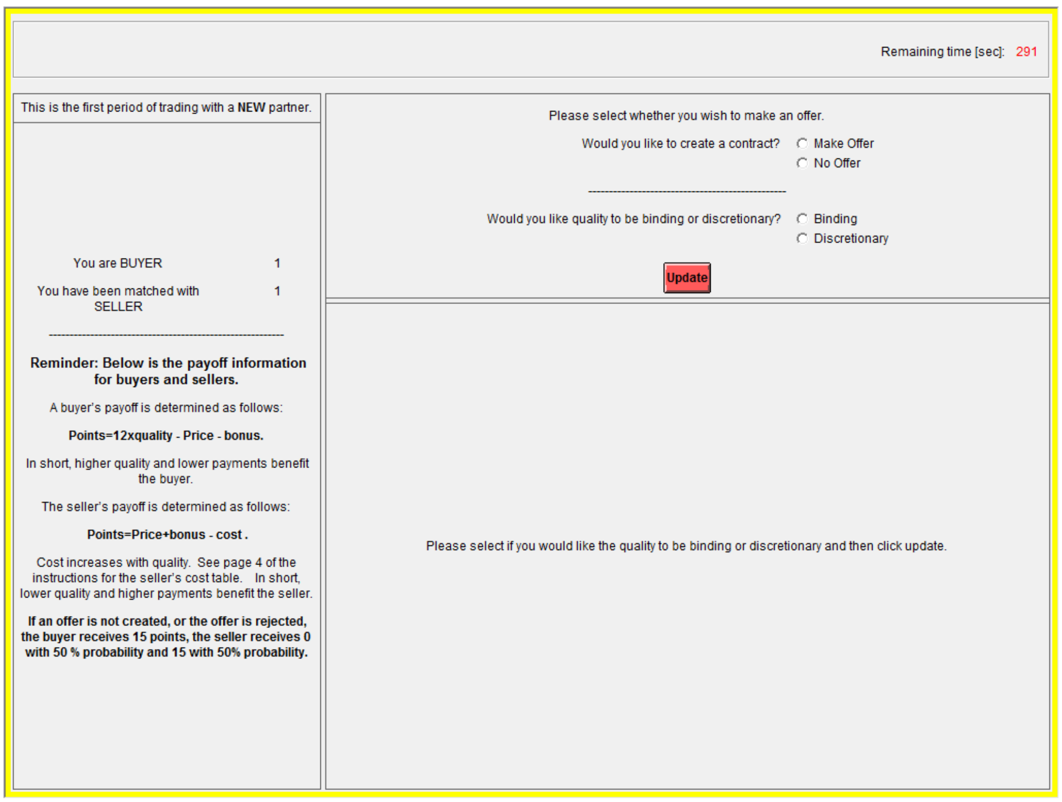

If buyer chooses “No offer,” and clicks “Update”, this is what s/he sees

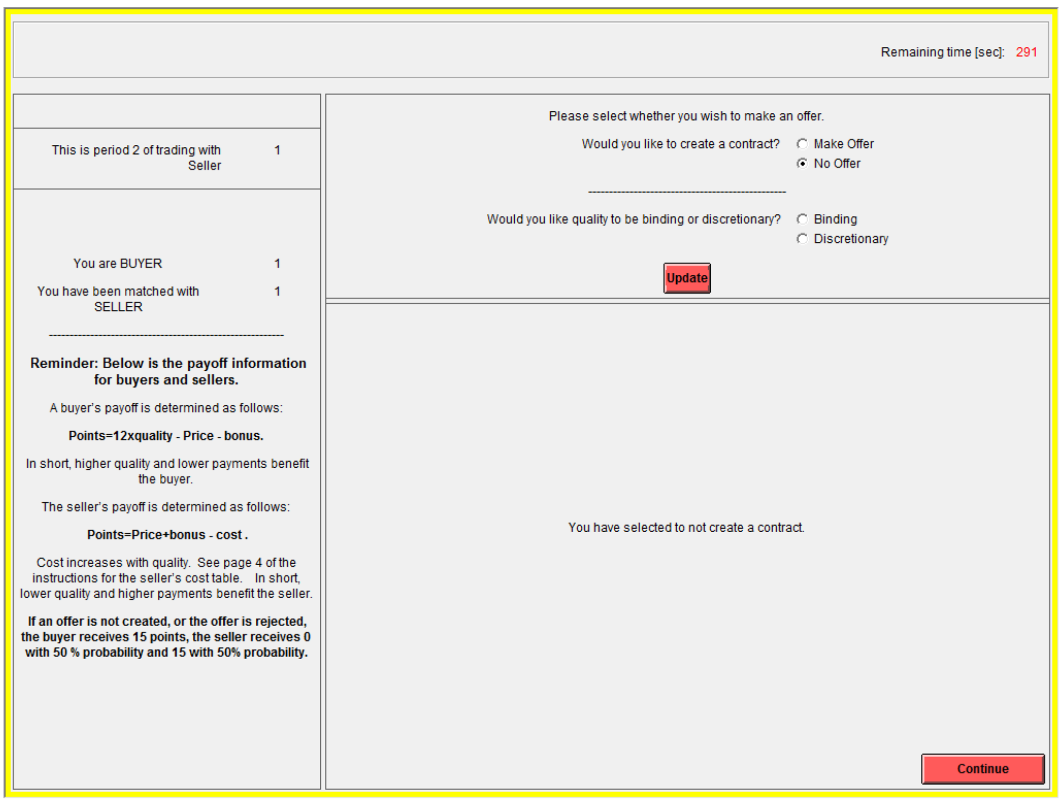

After pressing "Continue" on the previous screen, the subjects are shown the following end of period summary screen

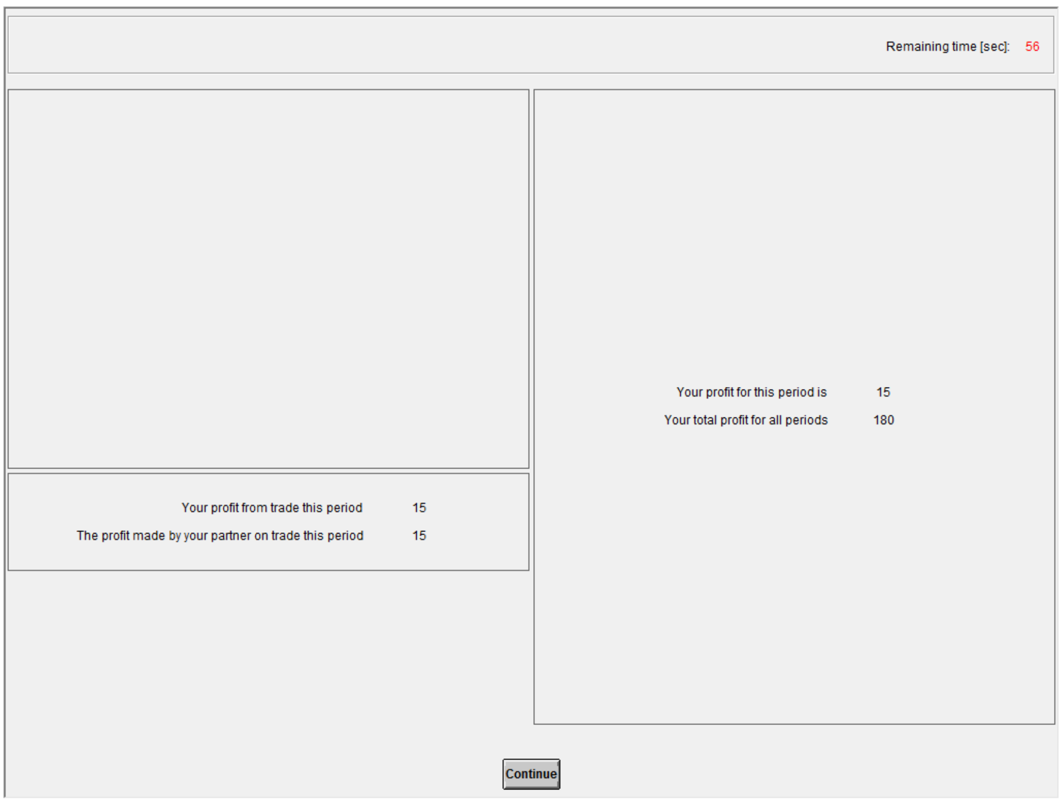

If instead the buyer had clicked “Make Offer” and “Binding” to create a binding contract that enforces quality and price, then the buyer offer screen (after clicking “Update”) changes to the screen below. Note the buyer must select the binding quality level (note: only 1 and 5 are verifiable qualities in the PE treatments) and enter an offered Price.

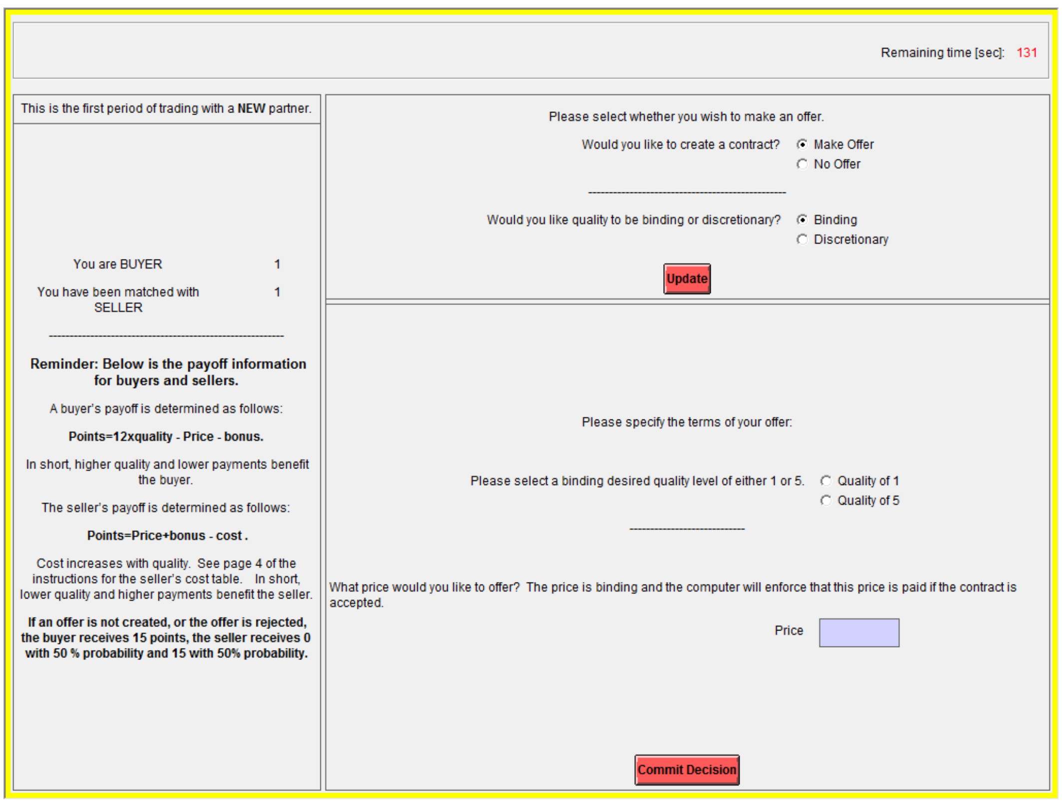

Suppose we enter binding quality of 5 and a price of 50. Then pressing "Commit Decision" takes us to the next screen for the buyer. The buyer waits at this screen because the seller must decide whether to accept or reject the contract. Note that the default bonus for a binding contract is 0 since the bonus plays no incentive role in a binding simple contract.

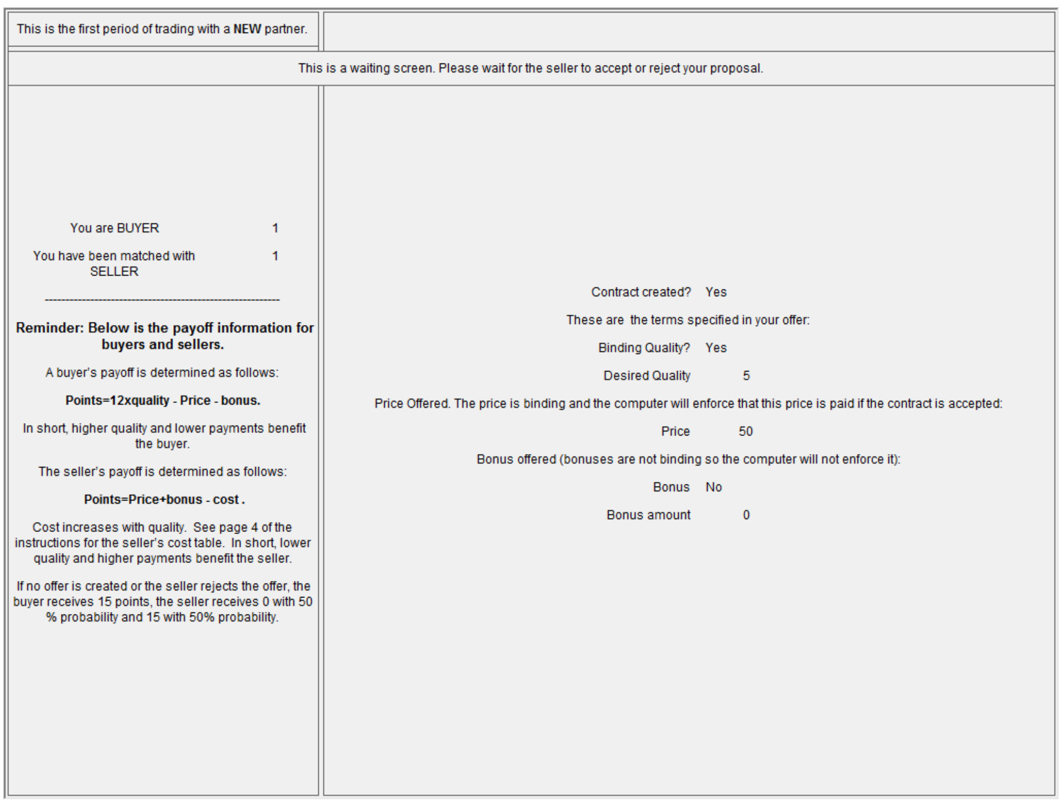

While the buyer is waiting, the seller sees the following screen.

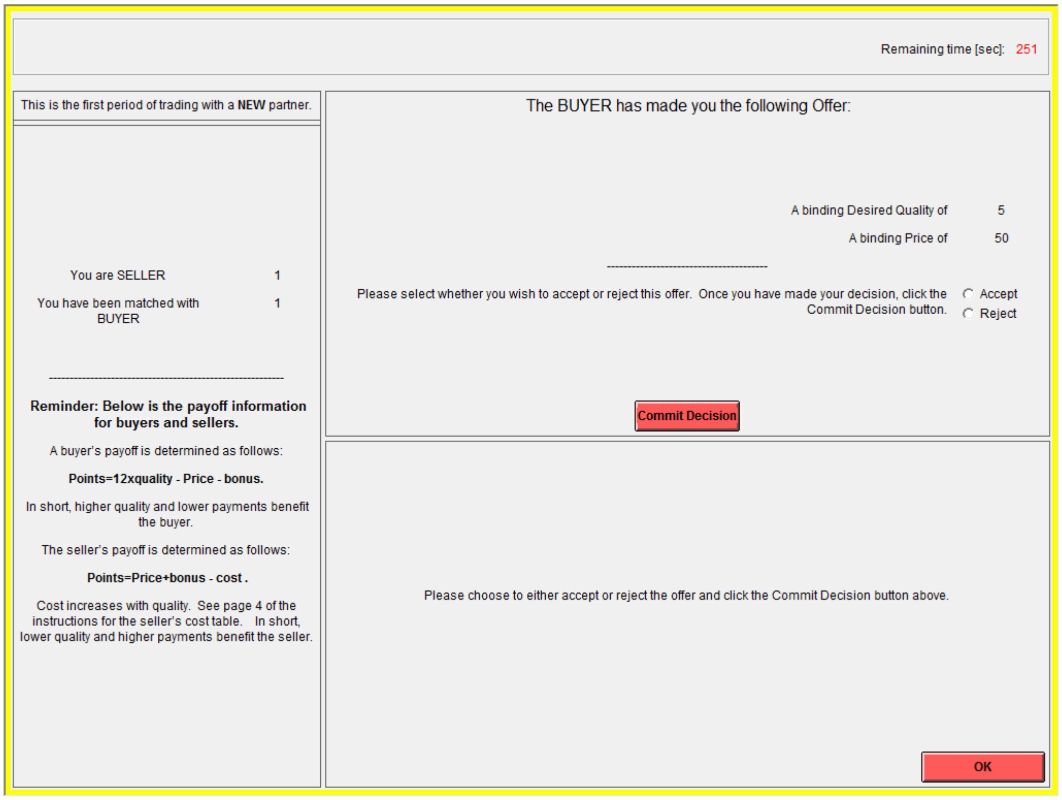

If the seller rejects the contract, then the seller is taken to the following screen (the buyer is shown an analogous screen)

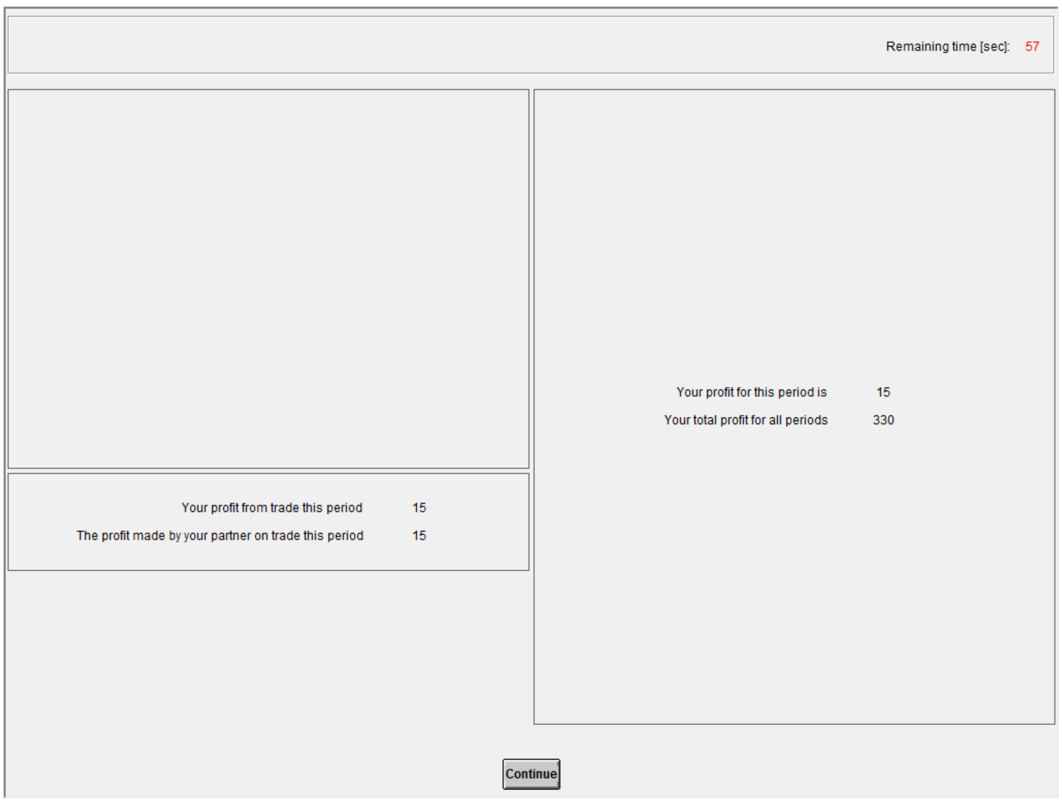

If the seller instead accepts the contract, then the trade is completed (there is no ex post discretion to choose quality or payments under a binding contract) and taken to the following screen (the buyer is shown an analogous screen)

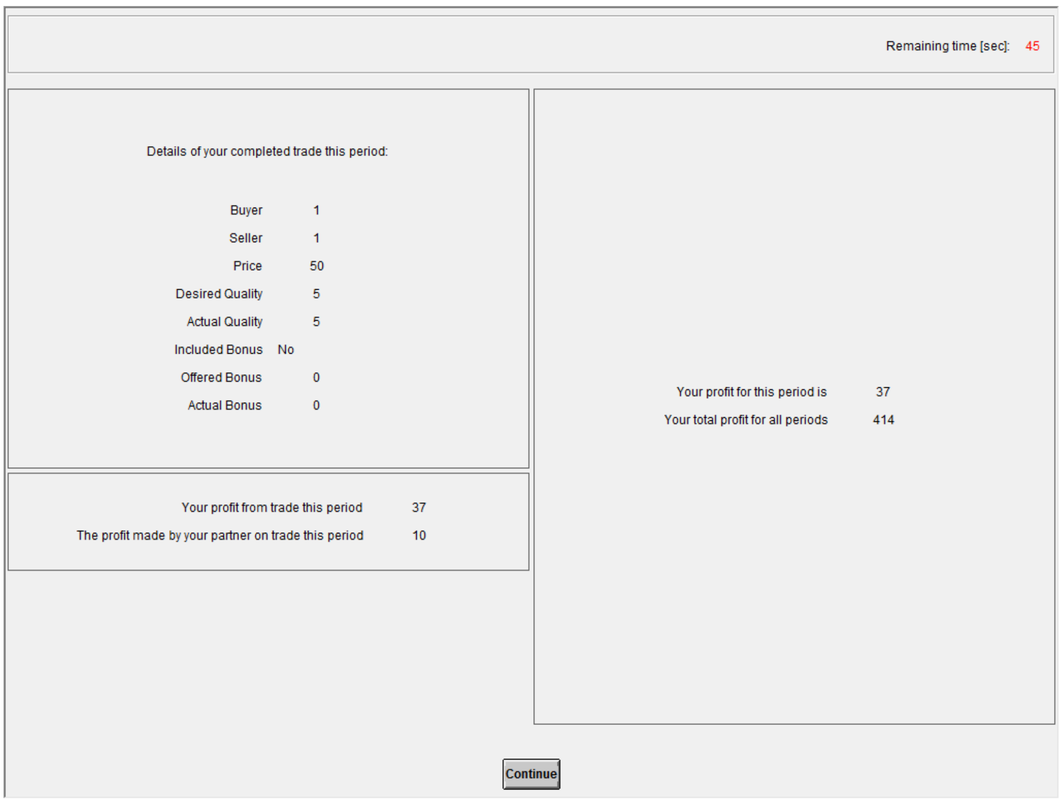

Now let's return to the buyer offer screen. Had the buyer chosen a discretionary contract, then the offer screen changes to the following:

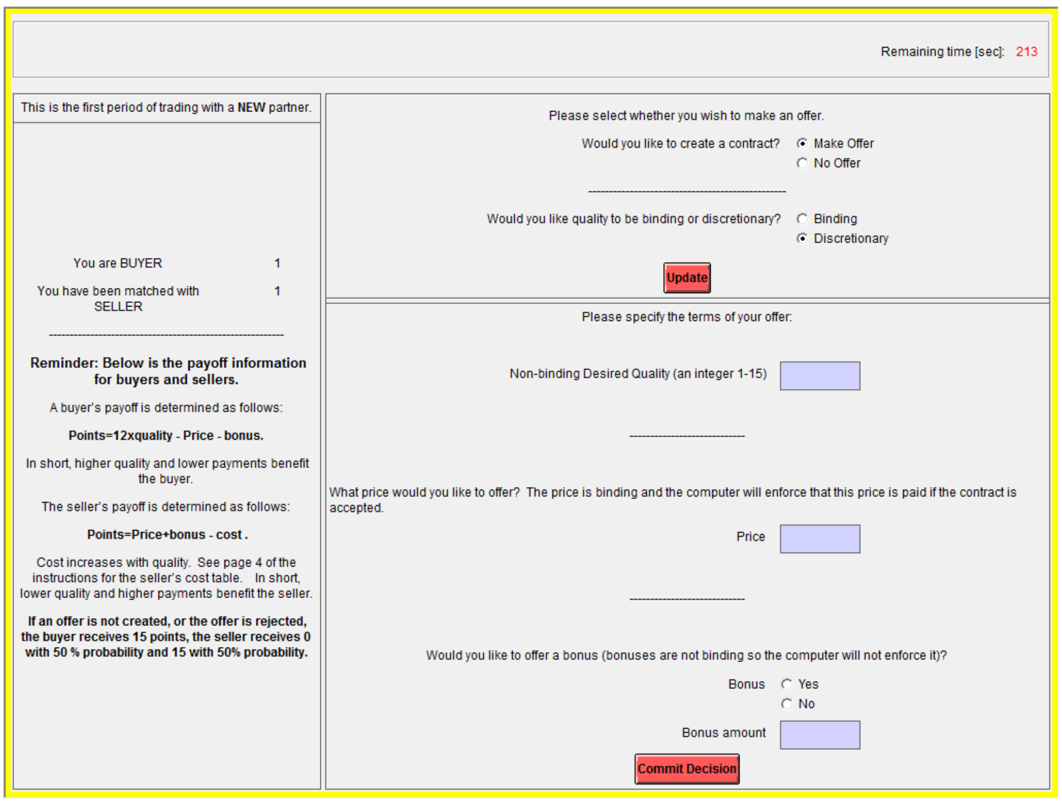

If the buyer offers a discretionary contract asking for Q=7, P=30 and B=30, then after clicking "Commit Decision" s/he is taken to the following waiting screen while the seller is making an accept or reject decision.

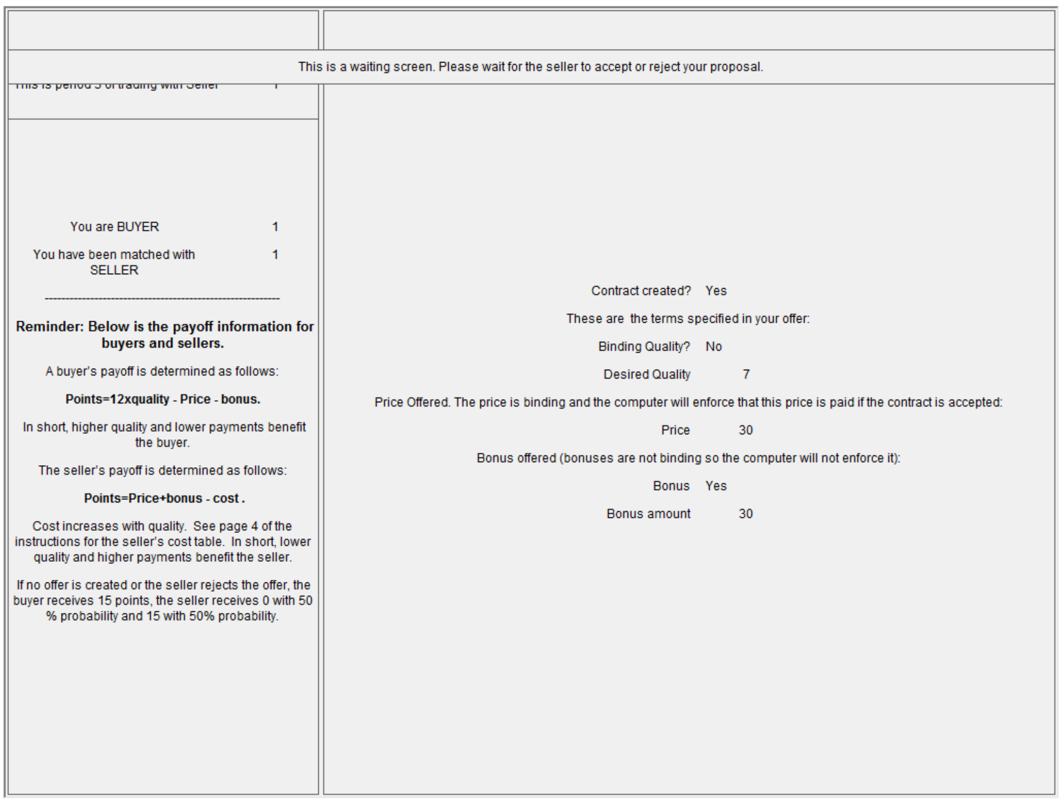

If the seller rejects the discretionary contract, then both buyer and seller are taken to the end of the period screen much like what has already been shown earlier. However, if the seller accepts the contract, her decision screen looks like the following (note: once s/he chooses accept, a quality determination box appears at the bottom of the screen):

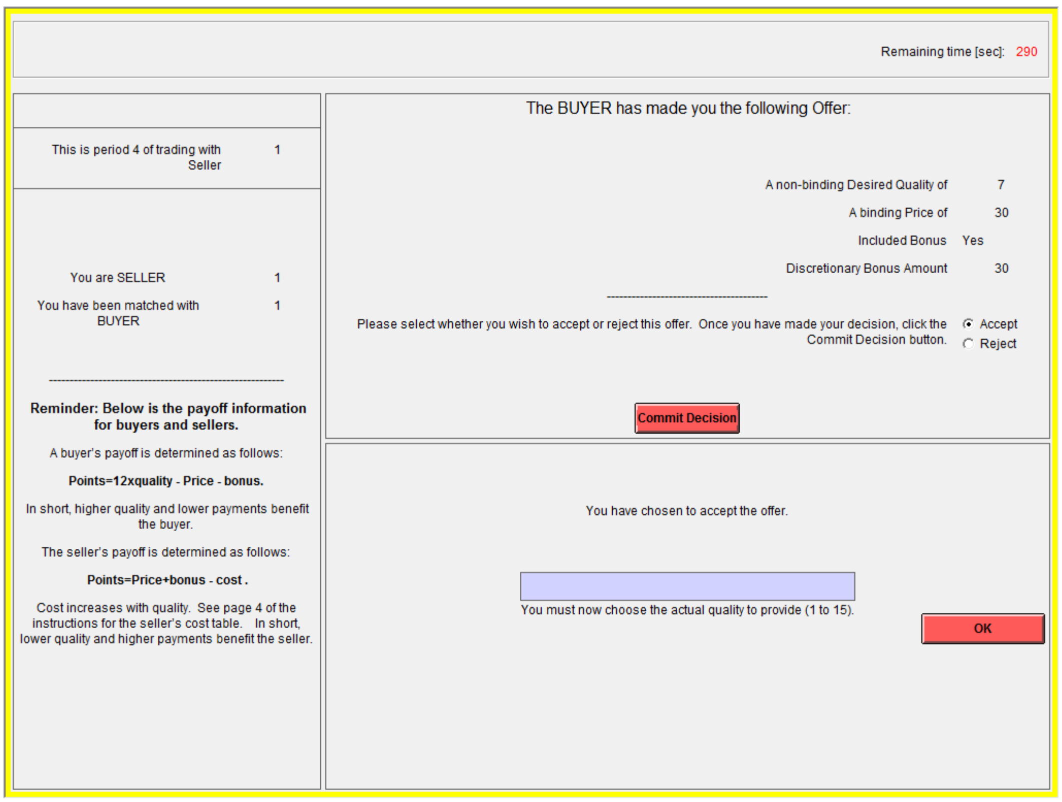

If the seller chooses an actual quality of q=5, s/he is taken to the following waiting screen.

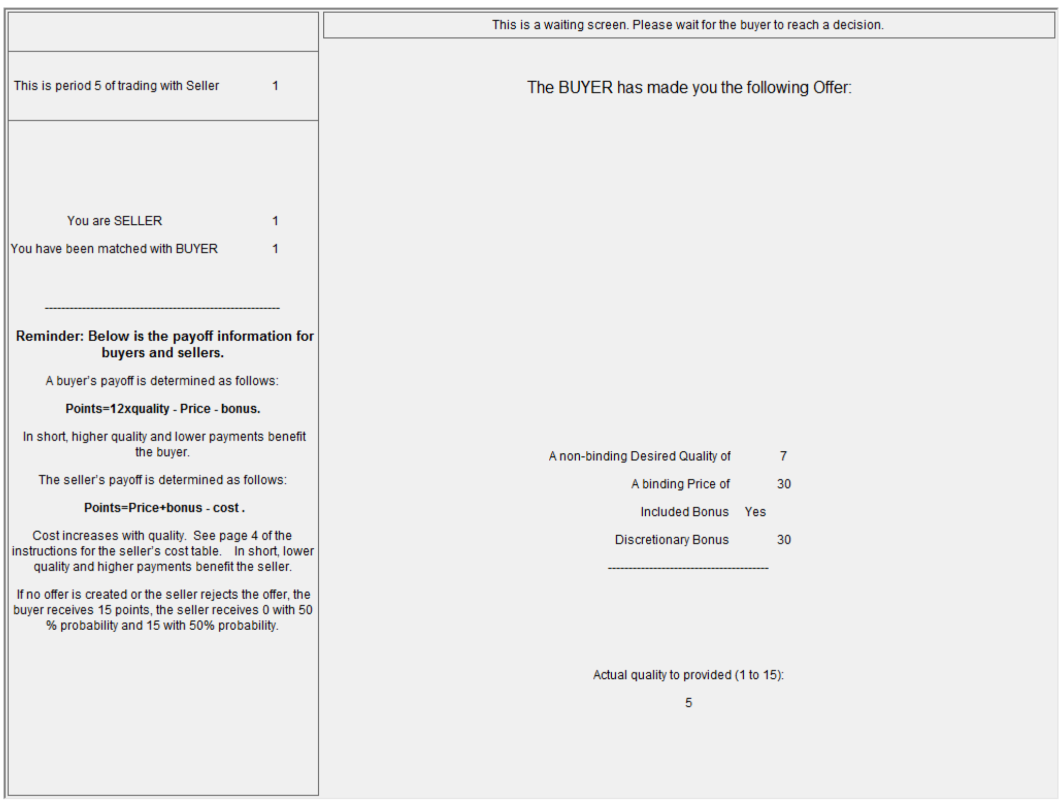

While the seller is waiting, the buyer is taken to the following bonus determination screen.

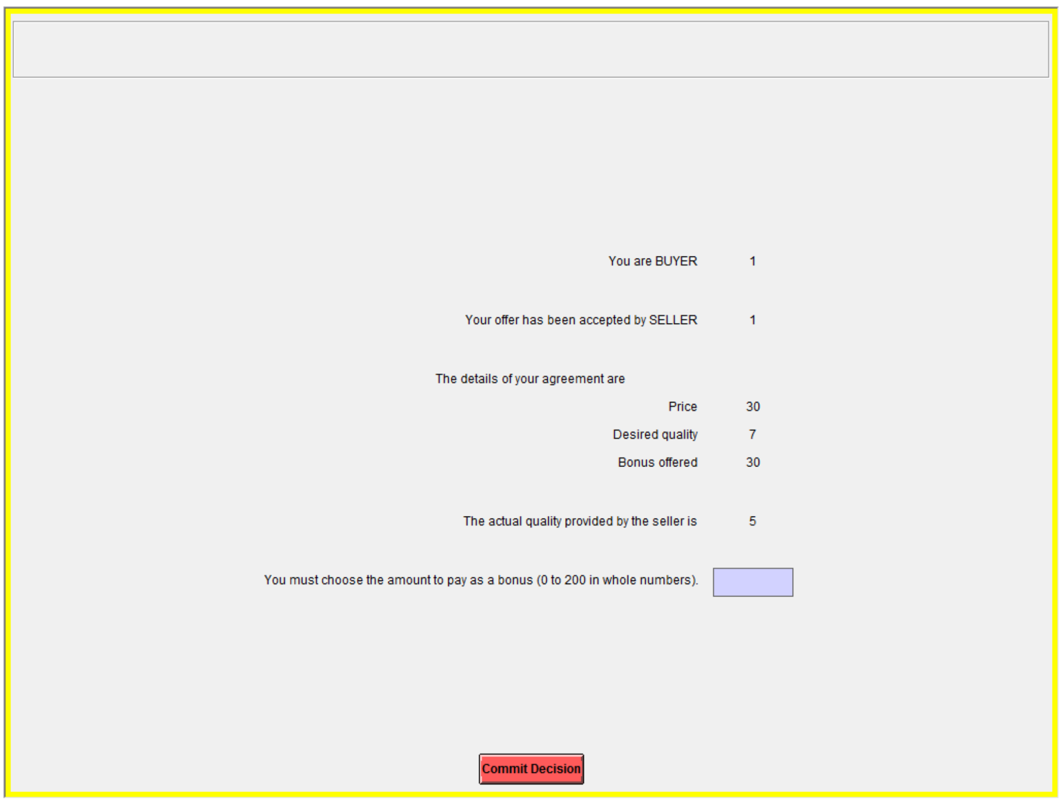

If the buyer pays an actual bonus of b=25 and then presses "Commit Decision," s/he is taken to the following end of the period summary screen. The seller sees an analogous screen.

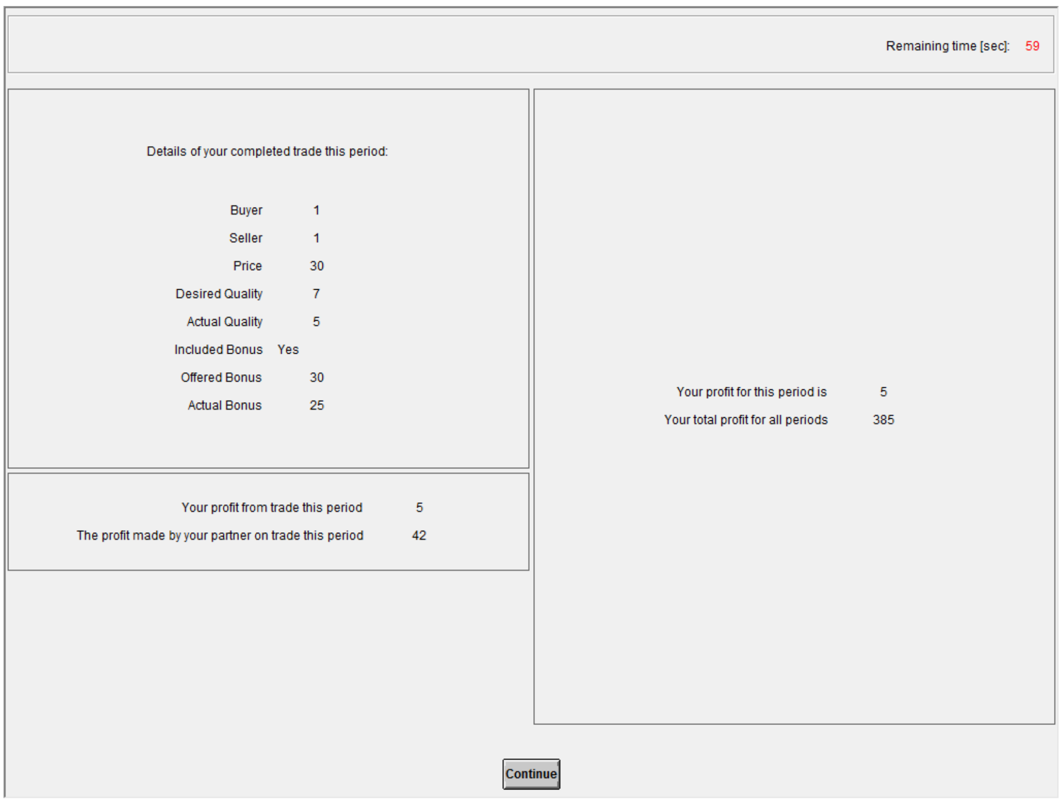

Once a period is over, both the buyer and seller see the following screen that determines their probability of trading with each other again the next period. A key point to note is that, as a practical matter, the realized draw of the continuation probability was simultaneously applied to all pairs of buyers and sellers in a session to facilitate orderly rematching when supergames terminate. In other words, either all pairs in the room continued or terminated in the same period. This made it easy to implement stranger matching. Nonetheless, to ensure saliency of the continuation probability, we forced each subject to press the "Reveal Draw" button to show them the realized draw (whether they will be rematched with the same partner or a new partner). To speed up the experiment, they were given a maximum of 15 seconds to press the button. After 15 seconds, the next period begins and the buyer offer screen appears. The experimenter announced to subjects whether they are rematched with the same person or matched with a new person. Moreover, the top left side of the decision screens for both the buyer and seller remind them how many periods they have been trading with the same partner. Thus, even if some subjects forgot to press the "Reveal Draw" button, subjects were still informed of the realized draw because we implemented multiple layers of prompts to ensure that subjects knew the draw.

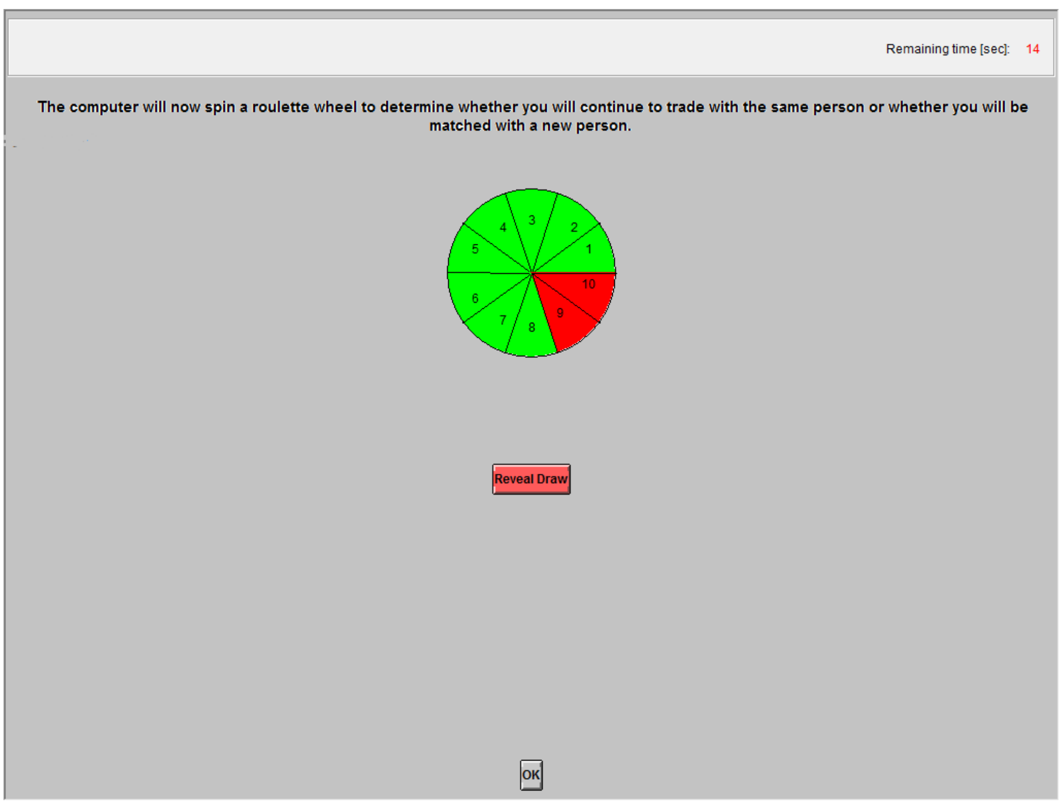

The next screen shows the revealed draw after a subject presses the "Reveal Draw" button