Experimental Economics manuscript No.
(will be inserted by the editor)

Appendix for "Putting Relational Contract Theory to the Test: Experimental Evidence"

Nisvan Erkal · Steven Y. Wu · Brian E. Roe

A Full Description of the Theoretical Model

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 theoretical results. We aim to provide a parsimonious unifying framework for many canonical results from the literature.

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 MacLeod and Malcomson (1989), 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 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 and non-contingent payment, p=P by default.

Let the principal's and agent's payoffs be $\pi = r(q) - p - b$ and $u = p + b - c(q)$ 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.

N. Erkal
Department of Economics, University of Melbourne, 111 Barry Street, Carlton, Victoria 3010, Australia
E-mail: n.erkal@unimelb.edu.au.

S.Y. Wu (Corresponding author)
Department of Agricultural Economics, Purdue University, Krannert Building, 403 West State Street, West Lafayette, IN 47907, U.S.A.
E-mail: sywu@purdue.edu.

B.E. Roe
Department of Agricultural, Environmental and Development Economics, Ohio State University, 225 Agricultural Administration Building, 2120 Fyffe Road, Columbus, Ohio 43210, U.S.A.
E-mail: roe.30@osu.edu.

A.1 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 it cannot detect more refined gradations in quality. Limited third-party verifiability allows for imperfections in performance measurement in the spirit of Baker, Gibbons and Murphy (1994), 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 enforceable to highly refined quality grades. Thus, our setup better matches stylized observations and allows us to nest both formal and informal contracts in a parsimonious framework.

Enforcement imperfections do not preclude the possibility of writing formal contracts, although imperfections do limit the set of available contracts. 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.

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, so 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 formal 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" formal 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 formal contract. Denote Q^f as the best contracted quality level.[1]

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

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

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

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 $Q^f = q^*$.

Similar to Baker, Gibbons and Murphy (1994), our model nests formal and informal contracts. Unlike Baker, Gibbons and Murphy (1994), we have a single performance measure rather than separately defining objective and subjective measures. This eases experimental implementation since subjects track fewer variables.

A.2 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 formal 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.[2]

In a relational contract, the stage game described above is infinitely repeated so that in each period t and for each history up to t, the parties follow 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 (optimal) contract is offered in every t.[3] 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:

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

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

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

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

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

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

[3] 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.

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 the principal's problem stated in* (A.2)-(A.6) *yields an optimal stationary contract that requests $\tilde{Q} \leq Q^*$ where Q^* is a request for first best quality. The associated 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:

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

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, suppose $\hat{Q} \geq Q^*$ 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, suppose that $\hat{Q} < Q^*$ so that $r'(\hat{Q}) > c'(\hat{Q})$. Then, the principal will contract for $\tilde{Q} = \hat{Q}$ and (A.7) binds with equality. We 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 $\frac{\delta}{1-\delta}\left[V(C) - \alpha\pi^f - (1-\alpha)\overline{\pi}\right] \geq B(\tilde{Q}) \geq \Big[c(\tilde{Q}) - c(\underline{q})\Big] - \frac{\delta}{1-\delta}\left[U(C) - \alpha u^f - (1-\alpha)\overline{u}\right]$. 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: $\hat{Q} < Q^*$: In this case, $r'(\hat{Q}) > c'(\hat{Q})$, 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 the first-best quality. The discretionary bonus simultaneously satisfies both the agent's and principal's SE constraints. The total promised payment satisfies both parties' IR constraints. This proposition directly leads to Hypothesis 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 informal 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. That is, suppose Q^f is the enforceable quality that yields the highest joint surplus among all contractible quality levels. A formal 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).[4] In short, an improvement in enforcement technology should cause some relational contracts to be replaced by formal 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 formal 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 formal 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 formal 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 the 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 $\tilde{u} = S(\tilde{Q}) - \tilde{\pi}$ 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 adhere 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 deviates, or by lowering P to yield per-period

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

payoffs $\tilde{\pi} = S(\tilde{Q}) - max\{u^f, \overline{u}\}$ and $\tilde{u} = max\{u^f, \overline{u}\}$ if the agent deviates. Such a contract continues to implement $\tilde{Q}$ because the self-enforcing conditions (part (ii) of Proposition 1) is independent of P. It 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., those that constitute a deviation), there is a family of relational contracts that implement $\tilde{Q}$ while delivering different payoff distributions. In this way, one can always construct an off-the-equilibrium-path contract that continues to implement $\tilde{Q}$, while holding the deviator to the payoff s/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 s/he would have been punished under the 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 Hypothesis 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 informal contracts are endogenously replaced with formal contracts.*

Proof First, note that $k = S(Q^*) - S(Q^f) = r(Q^*) - c(Q^*) - \overline{u} - \overline{\pi} - r(Q^f) + c(Q^f) + \overline{u} + \overline{\pi} = r(Q^*) - c(Q^*) - [r(Q^f) - c(Q^f)]$. 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 Bernheim and Whinston (1998) and to the substitutability between formal and informal contracts of Baker, Gibbons and Murphy (1994). Bernheim and Whinston (1998) 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 Hypothesis 4 in the main paper.

Another Bernheim and Whinston (1998) 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 want to have the discretion to adjust pay in response to the agent's action by utilizing a discretionary bonus contract. 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 Hypothesis 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 informal 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 informal contract if $S(\hat{Q}) > S(Q^f)$.*
3. *If $\delta < \underline{\delta}(\tilde{Q})$, then the principal switches to a formal 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 formal contract that implements Q^f if $S(\hat{Q}) > S(Q^f)$.

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

Corollary 3 leads to Hypothesis 3 in the main paper.

B Additional Analysis of the Potential Impact of Social Preferences and Semi-Grim Strategies on Efficient Punishment (Hypothesis 2)

This section extends the analysis of Hypothesis 2 regarding efficient punishment following a contract breach. We focus specifically on two potential explanations for why the results deviate from theory.

As mentioned in the paper, two potential explanations are offered for the deviation from theory. First, if subjects exhibit distributional social preferences, then extreme distributions such as those that call for one party to be held at her reservation payoff might not be feasible. Additionally, if subjects exhibit reciprocity, then there might be a tendency to excessively punish an uncooperative trading partner, even at a cost to oneself. Both of these could potentially lead to the breaking off of relational trading.

While our study was not designed to test for the impact of social preferences/reciprocity, the data in Table 1 can offer some clues to guide future research. It is important to note that for a contract to "qualify" to be included in the table, there must have been an executed trade in the prior period so that we can condition the contract on one of four possible trading outcome states. Thus, informal contracts for which there was no trade in the prior period, either because of contract rejection or because no contract was offered, were not included. An overwhelming majority of accepted contracts promised profits that were well above the reservation payoff of 15 despite the fact that buyers made take-it-or-leave-it offers. Promised profits are what the parties would earn if both parties adhere to the contract. The only extreme distributions observed came from the six accepted contracts in PE0.80 state (H,S) (which promised a seller profit of 1), the two in PE0.50 (H,S) (which promised a seller profit of -21), and the four in PE0.50 (S,H) (which promised a buyer profit of 11). What is interesting about these contracts is that they are not only extreme but excessively extreme in the sense that they promise payoffs to the deviating party that fall below reservation payoffs, which is what we might expect in the presence of negative reciprocity.

A second possible explanation of the deviation from theory is based on research by Breitmoser (2015) who shows that repeated prisoner's dilemma (PD) strategies are well described

Table 1: Accepted (Rejected) Relational Contract Terms After 1-Memory Histories

	(H,H)	(H,S)	(S,H)	(S,S)
PE80 treatment–Means				
B	54.48 (n/a)	12 (9)	24.14 (n/a)	35.71 (69.42)
P	36 (n/a)	42.5 (30)	48.71 (n/a)	42.55 (13.75)
Q	10.6 (n/a)	10.17 (10)	8.43 (n/a)	9.97 (9.83)
Promised Seller Profit	36.67 (n/a)	1 (-11)	35.57 (n/a)	25.71 (32.5)
Promised Buyer Profit	31.26 (n/a)	67.5 (81)	28.29 (n/a)	41.35 (34.83)
N	42 (0)	6 (2)	7 (0)	31 (12)
PE50 treatment– Means				
B	59.5 (20)	15 (n/a)	58.25 (15)	51 (98.71)
P	20 (45)	49 (n/a)	38.75 (3)	45.42 (9)
Q	9.75 (10)	13 (n/a)	9 (10)	10.29 (10.71)
Promised Seller Profit	31.75 (15)	-21 (n/a)	52 (-32)	38.79 (49)
Promised Buyer Profit	37.5 (55)	92 (n/a)	11 (102)	27.08 (20.86)
N	4 (1)	2 (0)	4 (1)	24 (7)

-Means for rejected contracts are in parentheses.

by 1-memory Markov semi-grim mixed strategies where parties cooperate with a high probability after mutual cooperation, defect with a high probability after mutual defection, and randomize with an intermediate probability when only one player has defected.

Returning to Table 5 in the main paper, we can see that the pattern of behavior is consistent with semi-grim strategies in the sense that continuation of relational contracting occurs with high probability after (H,H), but falls off if one or both parties shirk. The only difference between our results and those of Breitmoser (2015) is that play in our experiments appears to be "less-grim" after shirking has occurred. While Breitmoser (2015) finds that about 10% of the subjects 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 following (S,S). However, this is arguably to be expected in light of Hypothesis 2 where subjects can take advantage of the larger strategy space in a contracting game and make payment adjustments to reward and punish rather than simply end the relationship. Thus, subjects appear to move behaviorally in the direction of Hypothesis 2, but not to the extent predicted.

For further insights, Table 1 reports average contract terms across various 1-memory states for both accepted and rejected (in parentheses) relational contracts. One can see in Treatment PE0.80 that there is a clear pattern of contract term adjustments across 1-memory states. Using (H,H) as a benchmark, note that after (H,S), the promised profit level to shirking sellers drops dramatically. If instead, buyers shirk but sellers adhere (S,H),

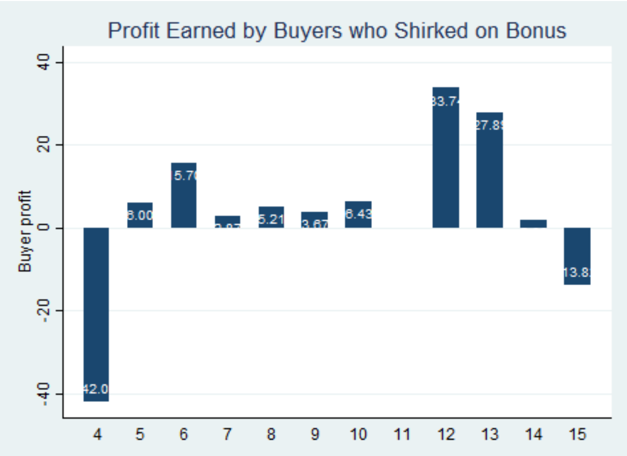

Fig. 1: *Profit earned by shirking buyers in PE0.50 across Q*

buyers do not seem to reward sellers with higher promised profits (35.57 vs 36.67), but they do offer higher fixed payments, P, and lower discretionary bonuses, B. This suggests that buyers try to reduce the strategic uncertainty faced by sellers. Finally, when both parties shirk (S,S), buyers respond by offering contracts that promise less profit to sellers (25.71 vs 36.67), but they also provide them with more security by raising P and lowering B. What is particularly interesting is that the rejected contracts (in parentheses) actually promise sellers a higher pay (32.5 vs 25.71), but they expose sellers to significantly more strategic uncertainty because P is significantly lower (13.75 vs 42.55) while B is significantly higher (69.42 vs 35.71).

C Additional Analysis of the Impact of δ on Q (Hypothesis 3)

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 $Q = 9.78$. We offer a couple of possible explanations and leave a more detailed analysis for future work.

First, because self-enforcement is so difficult in Treatment PE0.50, buyers may strategically design contracts for opportunistic purposes with no intention of self-enforcement. This conjecture is supported by the fact that shirking rates are so high among both buyers and sellers. 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 engaging in opportunism is to ask the seller to deliver a very high level of quality even if the buyer has no intention of paying the promised bonus. Figure 1 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 2 shows that binding contracts in treatment E implemented mean actual quality remarkably close to the first-best level of 12 (11.7 versus 12). Moreover, 48% of trades resulted in exactly the first-best quality. The few informal contracts used implemented $q = 7.14$ with only 5% implementing the first best.[5] Thus, when strategic uncertainty is eliminated, subjects chose values of Q that are remarkably close to the first best even though the first-best value of 12 was an interior solution and not an obvious focal point.

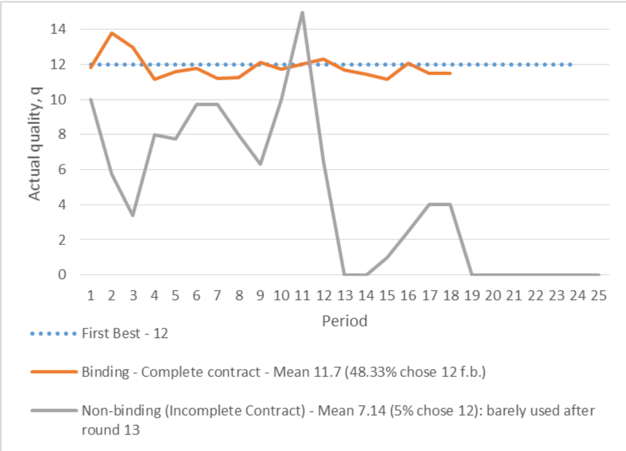

Fig. 2: *Actual q realized in Treatment E*

[5] Moreover, the informal contracts plot was volatile because very few trades used informal contracts. In many periods, only one or two trades were executed using informal contracts. In the later periods, many trades did not use informal contracts at all. These are the observations for which the plot touched zero quality.

D Instructions for Treatment PE0.80

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.

E Screen shots for Treatment PE0.80

This section contains the screen shots for Treatment PE0.80. The screen shots are presented in the same order as the sequence of moves within a stage-game.

Each period starts with the buyer offer screen:

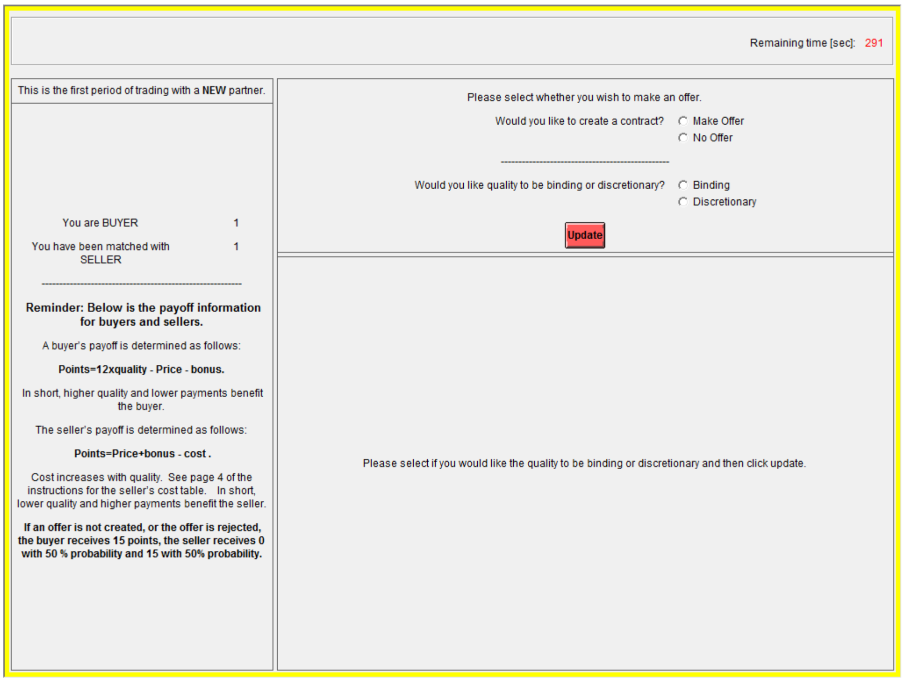

If the buyer chooses "No offer" and clicks "Update," this is what s/he sees:

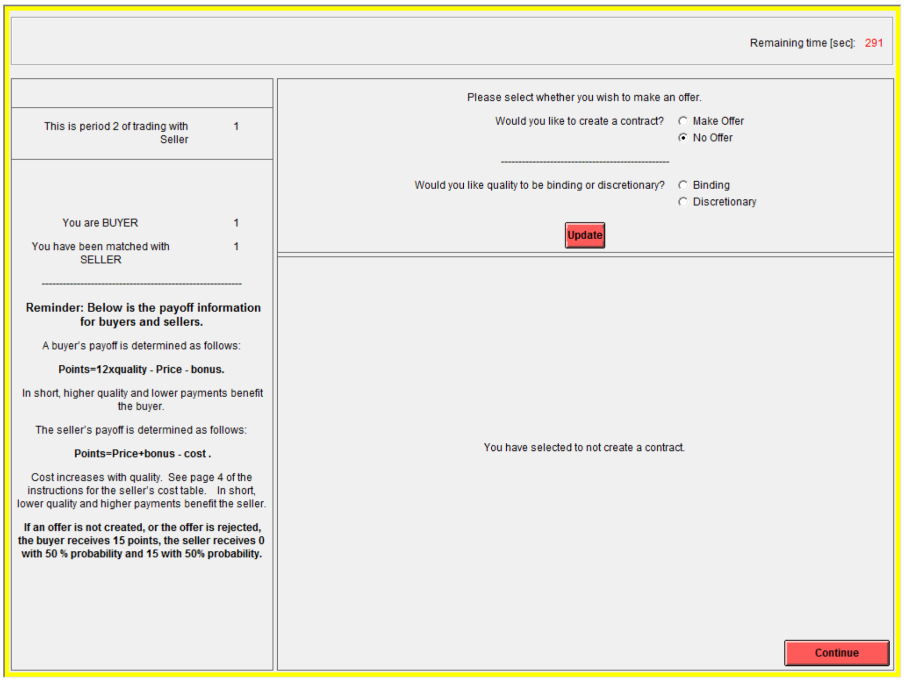

After pressing "Continue" on the previous screen, the subjects are shown the following end of period summary screen:

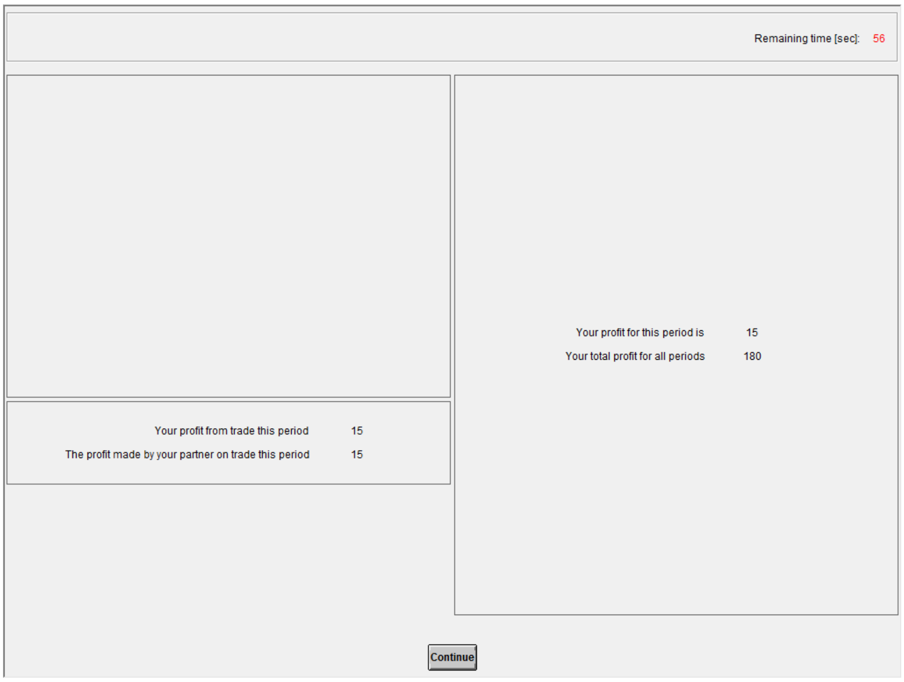

If instead the buyer clicks "Make Offer" and "Binding" to create a binding contract that enforces quality and price, then the buyer offer screen (after clicking "Update") changes to the screen below. The buyer must select the binding quality level and enter an offered price. Only 1 and 5 are verifiable qualities in the PE treatments.

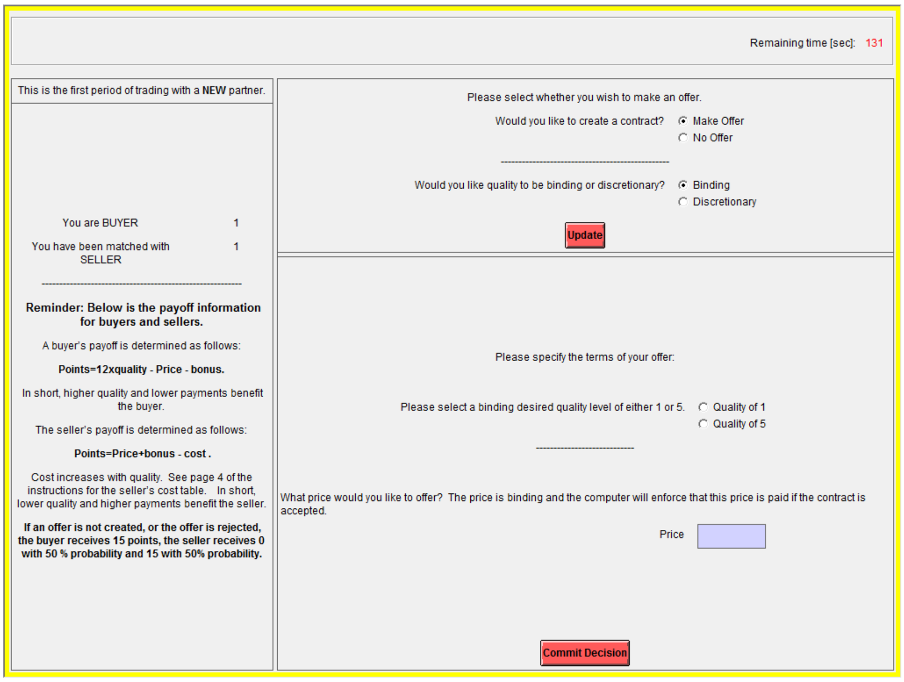

Suppose the buyer enters a binding quality of 5 and a price of 50. Then pressing "Commit Decision" takes us to the next screen for the buyer. The buyer waits at this screen because the seller must decide whether to accept or reject the contract. Note that the default bonus for a binding contract is 0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

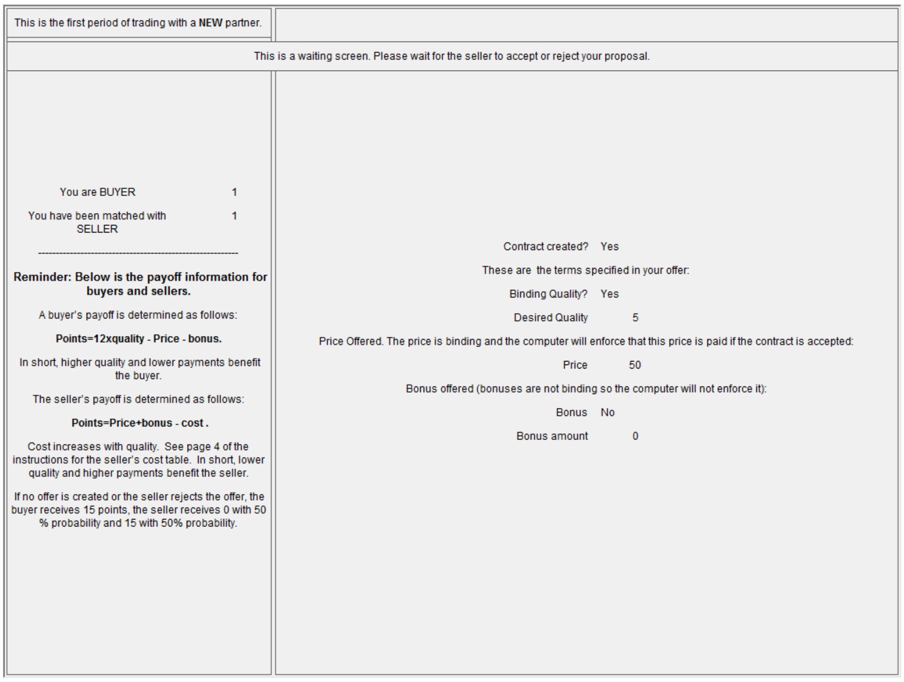

While the buyer is waiting, the seller sees the following screen.

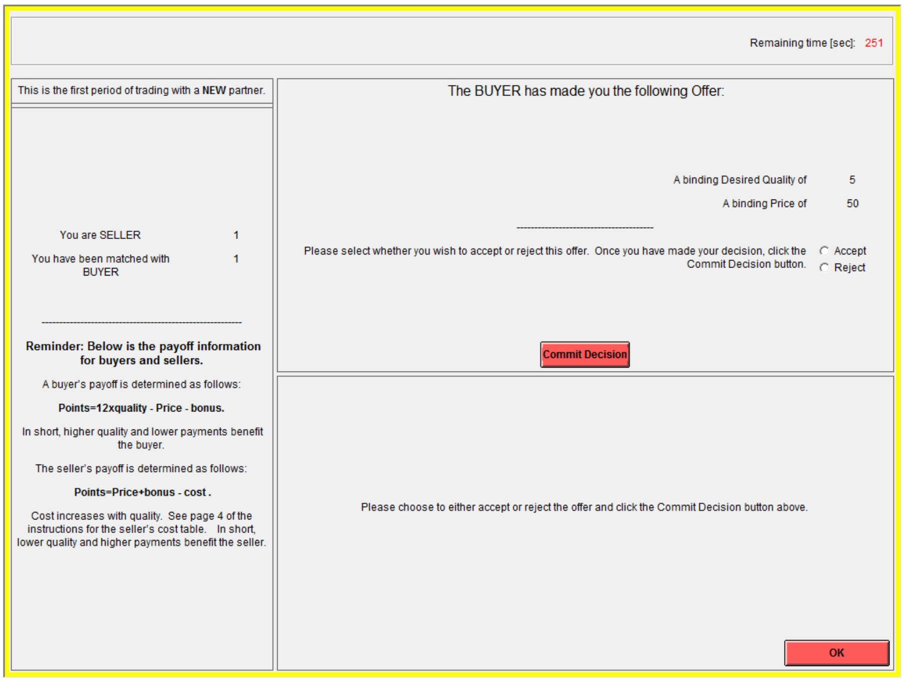

If the seller rejects the contract, then the seller is taken to the following screen. The buyer is shown an analogous screen.

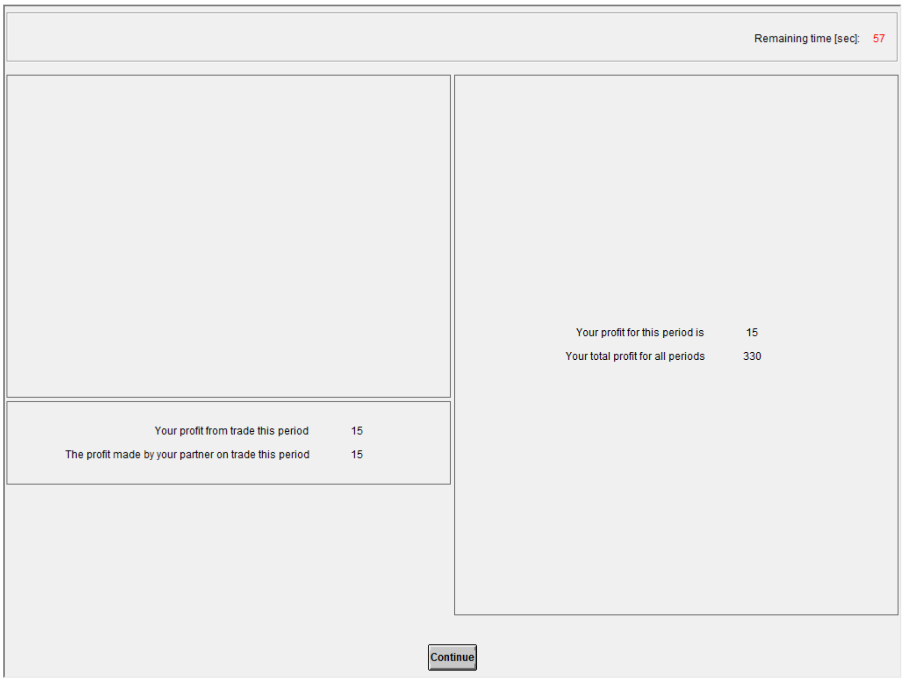

If the seller instead accepts the contract, then the trade is completed and the seller is taken to the following screen. (There is no ex post discretion to choose quality or payments under a binding contract.) The buyer is shown an analogous screen.

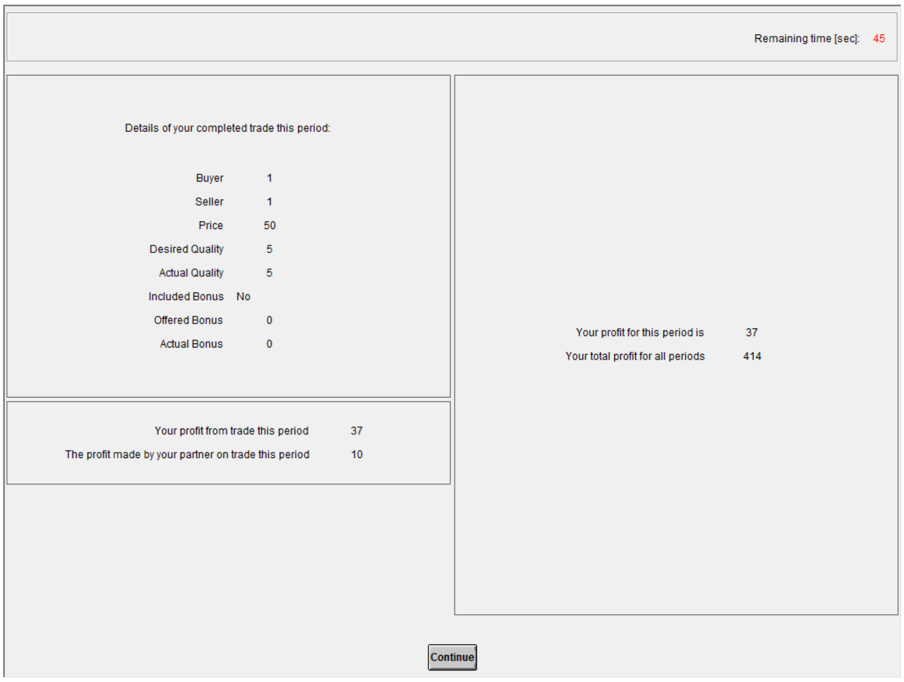

Now suppose the buyer chooses a discretionary contract. Then the offer screen changes to the following:

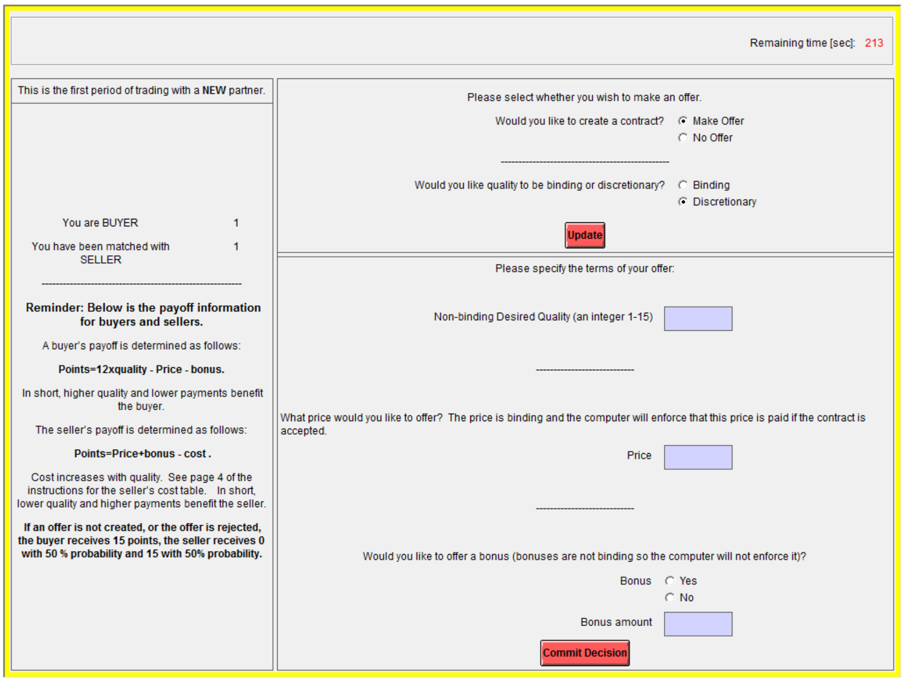

If the buyer offers a discretionary contract asking for Q=7, P=30 and B=30, then after clicking "Commit Decision" s/he is taken to the following waiting screen while the seller is making an accept or reject decision.

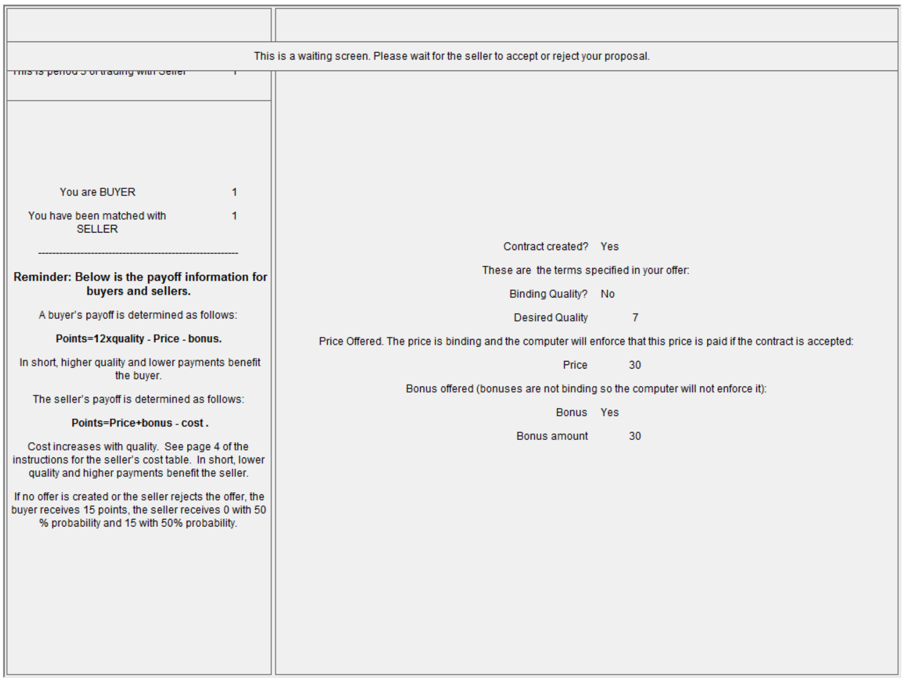

If the seller rejects the discretionary contract, then both the buyer and the seller are taken to the end of the period screen much like what has already been shown earlier. However, if the seller accepts the contract, the decision screen looks like the following. Note that once the seller chooses to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

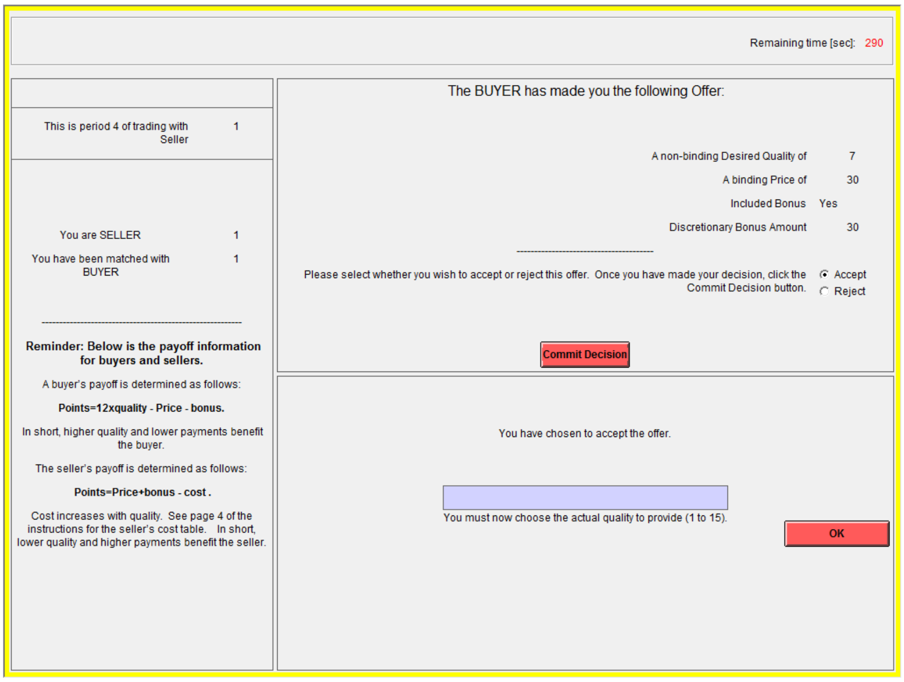

If the seller chooses an actual quality of q=5, s/he is taken to the following waiting screen:

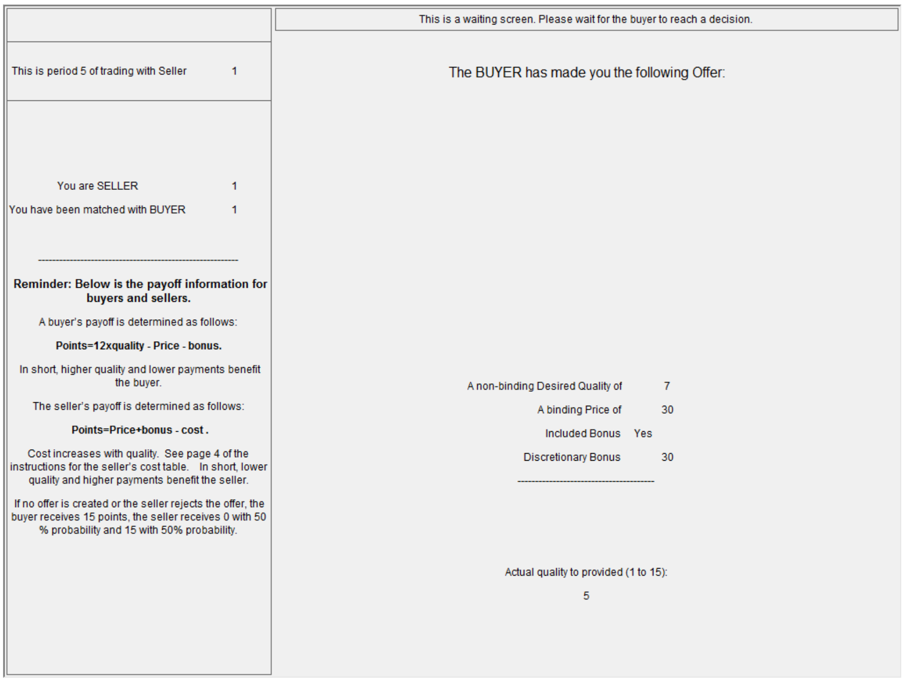

While the seller is waiting, the buyer is taken to the following bonus determination screen:

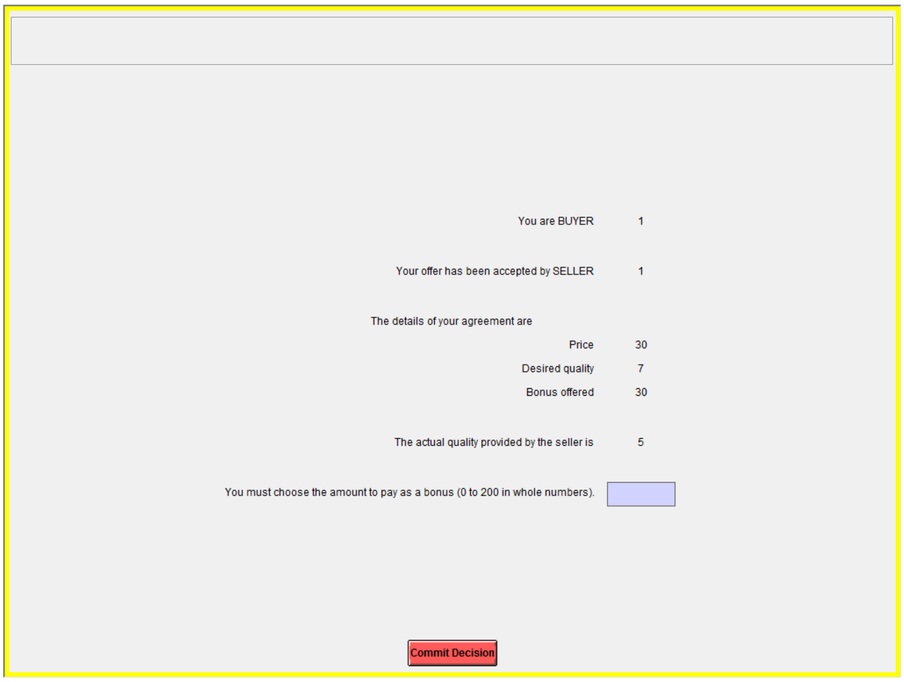

If the buyer pays an actual bonus of b=25 and then presses "Commit Decision," s/he is taken to the following end of the period summary screen. The seller sees an analogous screen.

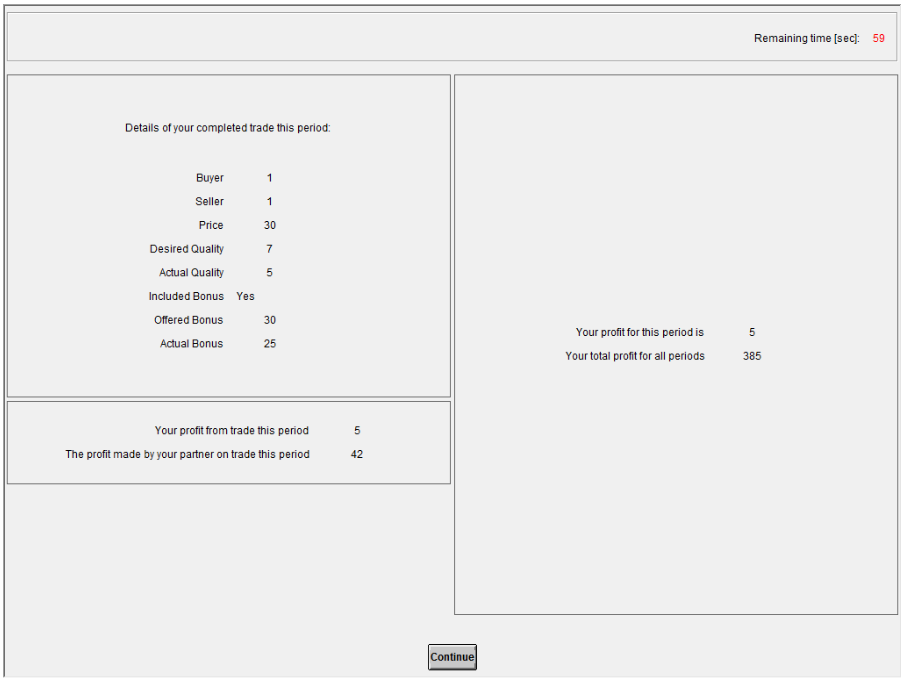

Once a period is over, both the buyer and the seller see the following screen which shows their probability of trading with each other again in the next period. A key point to note is that, as a practical matter, the realized draw of the continuation probability is simultaneously applied to all pairs of buyers and sellers in a session to facilitate orderly rematching when supergames terminate. In other words, either all pairs in the room continue or they all terminate in the same period. This made it easy to implement stranger matching. Nonetheless, to ensure saliency of the continuation probability, we asked 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). They are 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 announces whether subjects 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 the seller remind them of the number of periods they have been trading with the same partner. Thus, even if some subjects forget to press the "Reveal Draw" button, they are still informed of the realized draw because we implemented multiple layers of prompts to ensure that they are informed of the draw.

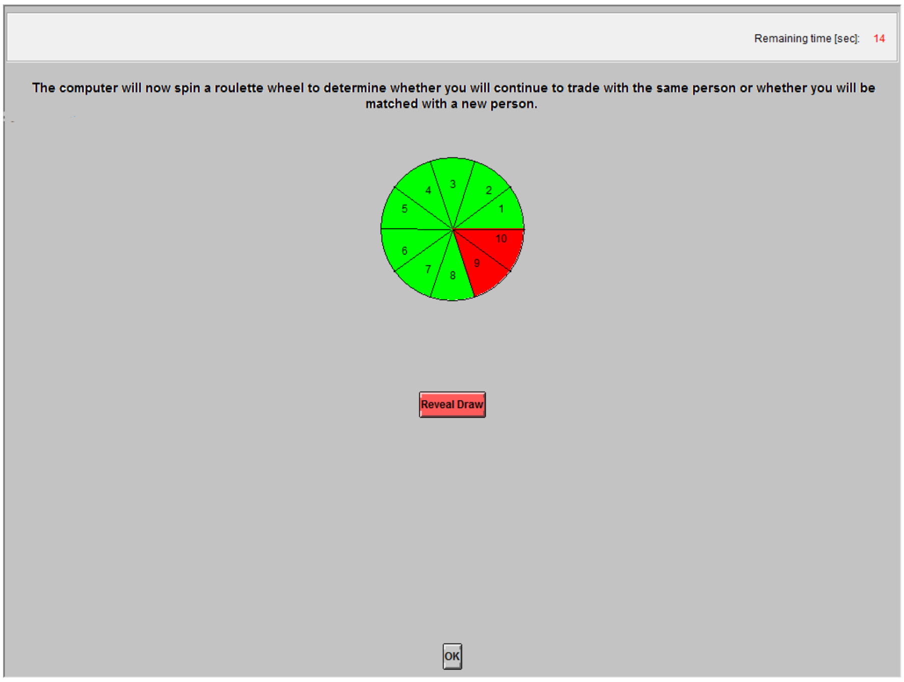

The next screen shows the revealed draw after a subject presses the "Reveal Draw" button.

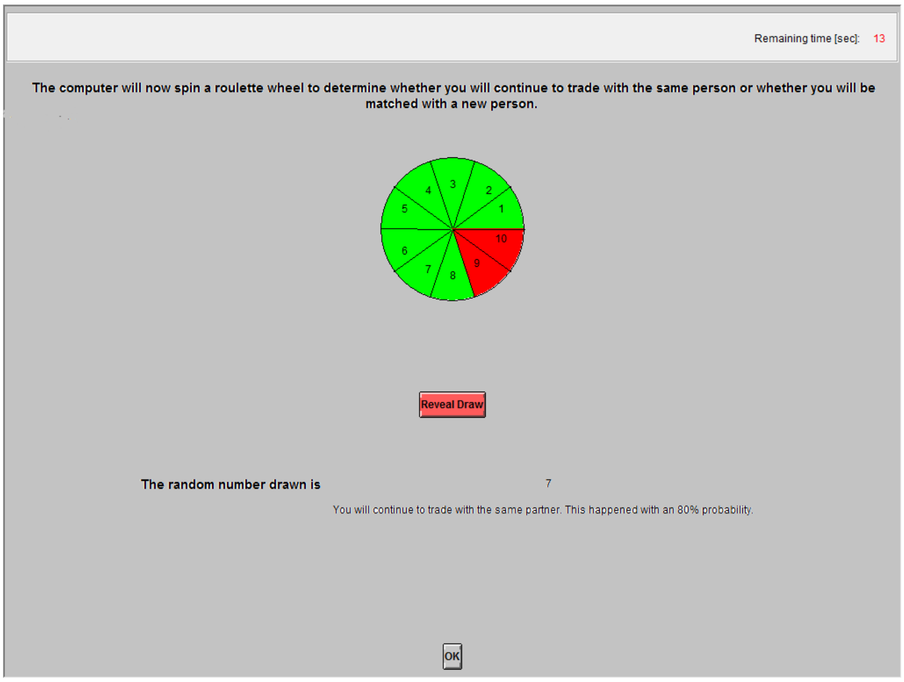

References

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.

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.

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.

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

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

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