8-26-2020

Trust and Trustworthiness in Procurement Contracts with Retainage

Matthew J. Walker

Elena Katok

Jason Shachat

Trust and Trustworthiness in Procurement Contracts with Retainage

Comments

ESI Working Paper 20-34

Trust and Trustworthiness in Procurement Contracts with Retainage*

Matthew J. Walker†, Elena Katok‡ and Jason Shachat§

August 26, 2020

Abstract

When product quality is unverifiable by third parties, enforceable contracts that condition price upon quality are not feasible. If higher quality is also costly to deliver, moral hazard by sellers flourishes, particularly when procurement is via a competitive auction process. Retainage is a contractual mechanism that presents a solution to the third-party unverifiability problem, by setting aside a portion of the purchase price. After delivery, the buyer has sole discretion over the amount of retainage money that is released to the seller. While generally a feasible contract form to implement, retainage introduces a moral hazard for the buyer. We use laboratory experiments to investigate how and when retainage might be successfully used to facilitate trust and trustworthiness in procurement contracts. We observe that retainage induces a significant improvement in product quality when there are some trustworthy buyers in the population, consistent with a model of fair payment norms that we develop. This improvement is realized at the cost of increased buyer-seller profit inequalities. We also observe that at high levels of retainage, there is a welfare-decreasing market unraveling in which sellers do not bid on contracts. Our results imply that retainage incentives can mitigate the tension between competition and cooperation arising from reverse auctions, but only at appropriate levels of retainage.

Keywords: trust, procurement, reverse auction, retainage, moral hazard
JEL Codes: C92, L15, D86

---

* The authors thank Daniel Zizzo and Ruth Beer for their constructive comments. In addition, the authors thank Timo Heinrich, Kenju Kamei, Nejat Anbarci, participants of the Utah Experimental Economics Conference 2020, Behavioral Operations Conference 2019, University of Texas at Dallas LBOE seminar series 2019, and Newcastle Economics Research Development Group for their helpful suggestions. Jennifer Burson provided valuable research assistance. This project received IRB approval from The University of Texas at Dallas Office of Research Integrity and Outreach (#MR19-015). Funding from the Economic and Social Research Council Award No. ES/R500963/1 is gratefully acknowledged.
† Durham University Business School, Durham, DH1 3LB, UK; email: matthew.j.walker@durham.ac.uk (corresponding author).
‡ Naveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080; email: ekatok@utdallas.edu.
§ Durham University Business School, Durham, DH1 3LB, UK; email: jason.shachat@durham.ac.uk.

1. Introduction

Technological advances allow enterprises access to increasing numbers of potential sellers. Reverse auctions leverage this increased access into greater competition and downward price pressure. In an interview with Fortune, General Electric's then CIO claimed to use reverse auctions for nearly one third of their total buy.[1] Regulations in the Unites States now also require the allocation of most public-sector procurement projects via competitive auction processes (Bajari et al. 2014). The implied cost-savings for buyers are in the region of 10 to 40% (Elmaghraby 2007). At a first glance, the use of price-based auctions would appear to improve procurement efficiency and benefit the buyer. Unlike in other market settings, however, in procurement, the auction often marks the beginning rather than the end of the transaction. For many goods, sellers determine their final product quality after the contractual allocation process.

As a consequence, price competition can erode trust in buyer-seller relationships. In complex projects, it is difficult for a buyer to write complete contracts that condition price on quality. Litigation is often expensive to pursue in practice. A classic moral hazard problem arises, in which a seller has limited financial incentive to restore product quality with costly action post-auction. For example, Emiliani and Stec (2005) observed that auctions encourage unethical cost-cutting behavior by sellers to preserve profit margins. This dynamic is likely to result in inefficiency. Thus, we see preferences among procurers for repeat purchases from a smaller group of certified suppliers and the establishment of reputational mechanisms.

In this paper, we explore an alternative solution to the seller moral hazard problem, by introducing price flexibility into the contract via *retainage*. A retainage provision, or retention as it is known outside the US, is a pre-agreed percentage of the contractual price withheld from a seller by the buyer. The buyer in this context might be a client, main contractor or sub-contractor withholding money from a lower tier. After product delivery, the buyer has discretion over how much of the retainage money to return to the seller. Retainage thereby permits the buyer to offer voluntary and unenforceable incentives to the seller for good performance (Raina and Tookey 2013). Whereas complete contracts require the capacity of third parties to verify and enforce quality-contingent prices, a retainage provision relies at least in part upon the buyer's trustworthiness to fairly compensate suppliers for quality delivered.

Retainage has its origins in nineteenth-century British railway construction, when it was set at 20% of contractual value (Bausman 2004). Today, typical provisions range from 3 to 10%. A construction sector consultation in 2017 by the United Kingdom Department of Business, Energy & Industrial Strategy (BEIS 2017) estimated the size of retainage monies held in England

[1] http://archive.fortune.com/magazines/fortune/fortune_archive/2008/07/21/105711270/index.htm?postversion=2008071010.

alone at £4.5 billion. Around 75% of contractors and 85% of clients surveyed had experienced retainage at some point during the preceding three years. Retainage provisions were included in the majority of ongoing construction contracts and most prevalent in contracts of high value.[2]

Where sellers produce a relationship-specific product prior to receiving compensation, they accept vulnerability to uncertain behavior during trade.[3] Retainage is due on *substantial completion* of a project, at which point half of the money is released immediately and the remainder is released after the expiration of a defects liability period. *Substantial completion* is difficult to verify, and generally is at the discretion of the buyer, creating conditions ripe for opportunistic buyer behavior. Measures to safeguard cash retainage vary by country.[4] Over half of contractors surveyed by BEIS had experienced full or partial non-repayment of retainage monies by buyers, who chose – often unjustifiably – not to return retainage monies due. This acted as a financial constraint and generated a counter-productive increase in procurement costs. Buyer recognition of the hidden cost of retainage has driven a downward trend in the maximum retainage provision permitted by several US states (ASA 2018). Thus, retainage introduces a new countervailing buyer moral hazard, shifting the burden of trust in procurement away from the buyer to the seller.

This buyer moral hazard problem is accentuated in many informal procurement arrangements, such as those based on a *letter of intent* (LOI). There are various types of LOI, from a mere handshake agreement stating the intention of parties to trade, to an interim contract that is – in theory – replaced by a binding contract on expiration, to a final but non-binding trade agreement in its own right. The use of LOIs for initiating works quickly and in the absence of a legally binding contract is "widespread" in the UK construction industry (Wevill 2015, p 29). If work proceeds based on an LOI, then the buyer, while not explicitly designating withheld monies as retainage, withholds a high percentage of the purchase price up front. It is not difficult to specify production cost structures which expose potential sellers to large losses and deters them from participating in requests for quotes. When sellers do not participate in the contracting process, we call this *market unraveling*.

---

[2] In the UK, standard building contracts are produced by the Joint Contracts Tribunal, which encourages the holding of retention monies until practical completion (for details, see https://www.jctltd.co.uk/).

[3] Özer and Zheng (2017) define trust in the supply chain as a voluntary decision "to accept vulnerability due to uncertain behavior of another (the trustee), based upon the expectation of a positive outcome" (p 495). Trustworthiness is a voluntary behavior "in a way not to take advantage of the trustor's vulnerable position when faced with a self-serving decision that conflicts with the trustor's objective" (p 497). The importance of the market environment in determining trust behaviors is documented by Özer et al. (2011).

[4] The European Commission (2009, clause 41) prescribes that retainage monies "are not paid until the satisfaction of conditions specified in the contract for the payment of such amounts or until defects have been rectified." In New Zealand, the 2015 Construction Contracts Amendment Bill provides additional protection for the payment of retainage monies to sub-contractors. In China, retainage applies to pre-specified defects liability periods and at the time of writing enjoy additional financial guarantees from the Agricultural Bank of China.

The peril of supplying under an LOI, without a concrete payment schedule, is demonstrated by a notable English contract law case. In 2005, RTS Systems won a competitive tender to supply improved food packaging for the German dairy manufacturer Müller. Work began based on an LOI and Müller paid RTS only 30% of the agreed price up front and a further 40% later on. After expiration of the LOI and repeated deferral in the execution of a binding contract, Müller alleged product defects and refused to pay RTS the remaining 30% of the tender price. A protracted and costly legal battle ensued, centered around the basis for which a contractual agreement existed. The Supreme Court Justice pronounced on judgement day that "the moral of the story is to agree first and to start work later" (RTS Flexible Systems v Molkerei Alois Müller [2010]). If LOIs are indeed widespread, they are liable to deter participation of potential sellers in the contracting process.

We use laboratory experiments to conduct the first controlled assessment of how retainage incentives influence procurement outcomes, while preserving access to a large pool of potential sellers and a commitment to procure at the lowest price.[5] By varying the percentage of the contract price withheld by the buyer until after delivery of the product, we shift the relative burdens of trust between buyer and seller. Specifically, we compare the performance of procurement contracts allocated at a price-based auction in which there is either (i) zero retainage – a fixed-price contract, (ii) a retainage provision set such that potential sellers can adjust their bid upwards to compensate for the increased risk of non-payment of monies, or (iii) a high retainage arrangement – analogous to an LOI – in which the buyer pays only a small percentage of the contract price up front, and consequently potential sellers cannot fully compensate for future production costs incurred by bidding higher. These three arrangements approximate the real-world market conditions discussed above. That is, there is either a one-sided seller moral hazard problem, a two-sided buyer and seller moral hazard problem, or a one-sided buyer moral hazard problem. Standard theory predicts that efficiency will be low in the first two contractual arrangements and nil in the third arrangement as the market unravels.

In our setup, payment of retainage monies ex post is voluntary and so, according to standard theory, generates no improvement in product quality. We develop an alternative model of *fair payment norms* to demonstrate that, if sellers believe there are some trustworthy buyers in the population, high product quality can emerge in anticipation of a reciprocal payment. By doing so, we contribute to a growing behavioral operations literature addressing how social preferences influence supply chain contracting (Beer et al. 2018, Cui et al. 2007, Davis and

---

[5] In a non-competitive setting, Hoppe and Schmitz (2011) observe in an experiment that renegotiable option contracts can help solve the hold-up problem if sellers are endowed with bargaining power. A similar framework has been applied to the supply chain setting by Davis and Leider (2018).

Hyndman 2018, Hu et al. 2017, Katok and Pavlov 2013, Loch and Wu 2008, Pavlov, Katok and Zhang 2020).

Our experimental results support the model of fair payment norms. In the baseline zero retainage arrangement, the market collapses to an inefficient, low quality outcome. This is consistent with prior laboratory studies of reverse auctions with seller moral hazard and zero price flexibility (Brosig-Koch and Heinrich 2014, Fugger et al. 2019). We find that a retainage provision can significantly improve product quality in a two-sided moral hazard environment. Yet a failure of potential sellers to fully adjust their bidding behavior in anticipation of variable retainage payments means that buyers appropriate the increased social surplus.[6] Sellers are vulnerable to loss when bearing the full burden of contractual trust in our high retainage arrangement. This often leads to market unraveling.

The present study builds on a well-established experimental economics literature examining trust and trustworthiness in principal-agent settings. Our baseline seller moral hazard environment integrates the gift-exchange game of Fehr et al. (1993) into an auction setting. In a typical gift-exchange game, participants are assigned to the role of either buyer or seller and participate in a two stage exchange. First, the buyer sets a price. Second, the seller produces a costly product quality and creates the trade surplus. In the absence of reputational considerations, if there is a preannounced and finite number of repetitions, then the seller should incur the minimum production cost possible and receive the lowest available price. In contrast, experiments of the gift exchange game without competition, typically observed a positive relationship between price and quality, yielding a Pareto improvement relative to the equilibrium prediction and lending support to Akerlof's (1982) gift-exchange hypothesis (e.g. Anderhub et al. 2002, Fehr et al. 1997, Fehr et al. 1998, Fehr et al. 2007).

The existence of positive and negative reciprocity is robust in the laboratory to the imposition of demanding market institutions (Fehr and Falk 1999). This behavior can be rationalized with theories of distributional social preferences (Bolton and Ockenfels 2000, Fehr and Schmidt 1999). Yet, during first price reverse auctions, in which buyers can only transact with the lowest-priced seller, competition forces bids down to minimum cost. High quality is then not a profitable seller strategy. The discontinuity in quality choice as a function of price is a direct result of the incomplete contracting model (Hart 1995, Hart and Moore 1988).

---

[6] Auction models in which contract renegotiation leads to lower prices ex post are considered by Waehrer (1995) and Wang (2000). Shachat and Tan (2015) consider an auction-bargaining mechanism in which sellers compete at an English auction to deliver an indivisible good to a buyer. They find that re-negotiated prices are below the winning bid although, contrary to the Nash prediction, final and initial prices are positively correlated in an experiment. Herweg and Schwarz (2018) investigate cost overruns and the possibility of price increases when the seller is endowed with bargaining power. Chang et al. (2016) find that wealth-constrained sellers can use the credible threat of default to obtain higher prices through renegotiation.

By designing different price-setting institutions, prior auction studies have shown that it is possible to improve trade efficiency in environments with seller moral hazard.[7] Most closely related to this paper are Fugger et al. (2019) and Brosig-Koch and Heinrich (2014). These studies restrict attention to the nature of the auction selection rule with fixed-price (i.e., zero retainage) contracts. Fugger et al. (2019) show that simply giving buyers the option to select a seller who did not place the lowest bid (called a buyer-determined auction) significantly raises prices and quality levels. Like in our study, interactions are one-shot, no reputation information is available and buyers must accept vulnerability to loss so as to incentivize high quality. The authors employ a multi-level cost and quality design, across two different buyer valuation schedules. Their experimental data reveal that buyer-determined auctions yield a robust improvement in cooperation and efficiency and they rationalize this finding using a model of inequity-averse preferences.

Brosig-Koch and Heinrich (2014) consider a similar buyer-determined auction setup in which buyers can condition procurement acceptance decisions on past seller performance. Providing reputation information in this way significantly increases buyer profits and procurement quality, relative to a binding price-based auction format. Unlike in Brosig-Koch and Heinrich (2014) and Fugger et al. (2019), we deliberately consider a binding price-based auction format to rule out the possibility that buyers choose the bid as a signal of quality. Buyer-determined auctions and reputational systems clearly have an important trust-building role in procurement. By contrast, we are interested in isolating the impact of contractual incentives during price competition without the confounding effect of alternative selection mechanisms.

The rest of the paper proceeds as follows. In section 2, we outline the setting and derive competing analytical results from the standard theory and a model of fair payment norms. In section 3, we derive testable hypotheses and summarize the experimental design. In section 4, we present our experiment results and conduct a formal statistical analysis. In section 5, we estimate parameters of our fairness model from the data and benchmark these with previous literature. In section 6, we conclude by drawing implications for managerial decision-making and procurement contract design.

[7] The first experimental analysis of procurement contracts in reverse auctions with moral hazard was conducted by Cox et al. (1996). They compared fixed-price and cost-sharing contracts. Their main finding was a trade-off between budgetary expense and efficiency. Although contracts with a greater cost-sharing element involved less procurement expense, they were also less efficient due to heightened seller moral hazard. Cost-sharing arrangements tend to be more appropriate for complex projects, accompanied by low degrees of design completeness (Bajari and Tadelis 2001). For the procurement of complex goods, a buyer commitment to negotiate with one seller can outperform an auction when there is adverse selection, scope for product improvements and/or costly renegotiation (Herweg and Schmidt 2017).

2. Model and Theory

Consider a one-shot procurement interaction in which one buyer seeks to purchase one unit of an indivisible product from a group of n potential sellers, indexed by i. A first price sealed bid auction determines selection of the winning seller and the contract price. Auction participation is voluntary for sellers, and the buyer can choose to not purchase after observing the contract price.[8] The winning seller can produce either a high- or low-quality unit, but the setting prohibits quality contingent contracts.

In this setting, the contract price is less binding than usual as purchases are made with retainage provisions. Such a provision includes a retainage proportion ρ, which is a fraction of the contract price withheld from the seller until after production.[9] A general interpretation of ρ is the degree of contractual flexibility. After production of the unit, the buyer learns of the chosen product quality. The amount of retainage released to the seller is then at the buyer's discretion. The retainage proportion thus regulates each party's trust burden: at low levels, the buyer possesses limited insurance against low quality production; at high levels, the seller is vulnerable to financially damaging retainage return decisions.

Figure 1 displays the sequence of events. In the Bidding Stage, each potential seller simultaneously submits his bid, b_i, or chooses not to participate. If at least one seller submits a bid, the one submitting the lowest bid wins the auction and the contract price is the winner's bid, $p = min\{b_1, \cdots, b_n\}$. Ties are broken randomly. If no bid is submitted, all parties earn zero and we call this outcome "market unraveling". When an auction succeeds, the profile of bidding-stage actions is announced before the next stage.

In the Procurement Stage, the buyer either rejects the auction outcome, resulting in all parties earning zero, or he accepts the outcome and pays $(1 - \rho)p$ to the winner. We call a buyer rejection "transaction failure". The buyer's Procurement Stage action is $a \in \{a^0, a^1\}$, where a^0 is a rejection and a^1 is an acceptance.

In the Production Stage, the seller chooses to produce either a high- or low-quality product, $q_i \in \{q^L, q^H\}$. The seller incurs a sunk production cost $c(q^j)$ for quality level j. Production cost schedules are the same across potential sellers and this is common knowledge. A seller's cost of producing high quality is strictly greater than his cost of producing low quality, $c(q^H) > c(q^L)$. Likewise, the buyer's valuation of the unit is increasing in quality and given by the expression $v(q^j)$. We denote $s(q^j) = v(q^j) - c(q^j)$ as the surplus created by the transaction.

---

[8] The inclusion of a participation decision for both buyer and seller ensures trust is a voluntary decision.
[9] When $\rho = 1$, bidding simply constitutes a promise of payment, i.e., uninformative cheap talk.

Figure 1. The sequence of events in our procurement model.

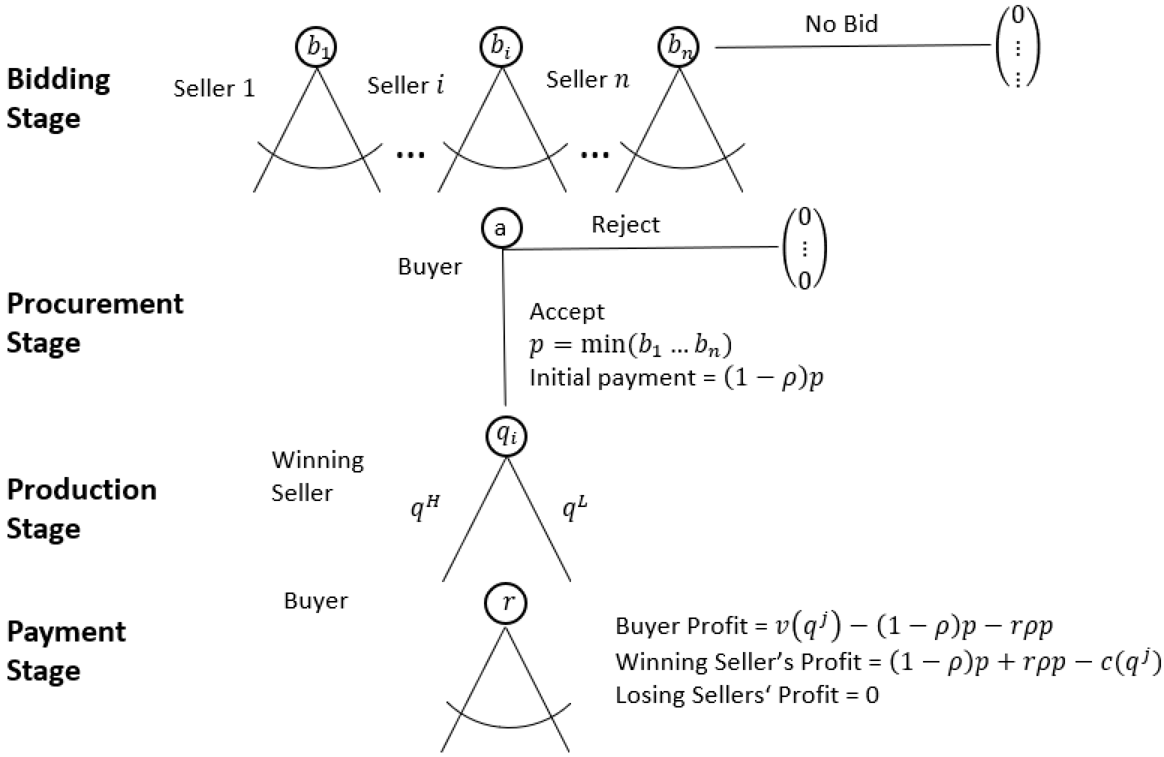

Notes: This is an extensive game tree representation of the strategic interaction. Sellers move first and either submit a bid b_i at auction or choose not to participate. If a market forms, the buyer can either accept to trade with the lowest bidder and make a guaranteed payment equal to $(1-\rho)p$, or refuse the transaction. The winning seller then selects to produce a high quality, q^H, or low quality, q^L, product and incurs the production cost $c(q^j)$. The buyer is informed about the product value, $v(q^j)$, and decides on a discretionary proportion r of the retainage money ρp to pay to the winning seller. The winning seller earns a profit equal to the difference between total payment received and the cost incurred. The buyer earns a profit equal to the difference between value received and the total payment made. Non-trading parties earn zero profit.

From a welfare perspective, trade is preferred to no trade and surplus is increasing in quality, i.e., $s(q^H) > s(q^L) \geq 0$. Finally, in the Payment Stage, the buyer observes the quality and then selects a fraction, *r,* of the retainage money to return to the seller.

Formally, a potential seller *i*'s strategy has two components. These are a Bidding Stage action $b_i \in \{[c(q^L), v(q^H)], b^0\}$, where bids can be submitted from a continuous interval between the seller's minimum production cost and the buyer's maximum valuation for the unit (inclusive), and b^0 is non-participation; and a quality choice function $q_i(b_i | b_i = p, a = a^1)$. A buyer's strategy also has two components: a procurement decision function, $a(p)$; and a retainage payment function, $r(p, q_i |\, a = a^1)$.

The profits of transacting buyers and the winning seller are,

Buyer's profit : $\quad \pi_B = v(q^j) - (1-\rho)p - r\rho p$

Seller's profit : $\pi_S = (1-\rho)p + r\rho p - c(q^j)$

Our primary interest is in whether and under what conditions there exists an equilibrium solution that implements high quality. High quality is implementable if it yields a non-negative expected payoff to the transacting parties (buyer and seller participate voluntarily) and it is incentive compatible (preferred to low quality by the seller).

2.1 Standard Theory

In this first approach, we assume the buyer and sellers are expected profit-maximizers. We proceed to solve for the subgame perfect equilibria using backward induction.[10]

Proposition 1.

- **A.** *For $\rho \in [0, 1 - c(q^L)/v(q^H)]$, there is a unique subgame perfect equilibrium in which $b_i^* = c(q^L)/(1-\rho)$ for all i, $a^* = a^1$, $q_i^* = q^L$ and $r^* = 0$.*
- **B.** *For $\rho \in (1 - c(q^L)/v(q^H), 1]$, there is a unique subgame perfect equilibrium in which $b_i^* = b^0$ for all i, and the market unravels; off the equilibrium path, $a^* = a^1$, $q_i^* = q^L$ and $r^* = 0$.*

Proof. **A.** The trade quality and contract price are known in the Payment Stage. Since the buyer's profit is decreasing in the retainage return fraction r, the buyer will always choose $r = 0$.[11] In the Production Stage, with certain zero retainage payment, low quality always yields higher profit than high quality. In the Procurement Stage, the buyer anticipates this and accepts to procure for all $p \leq v(q^L)/(1-\rho)$. In the Bidding Stage, each seller's expected profit maximizing bid is the cost of low quality marked up in proportion to the retainage money vulnerable to loss, generating an expected profit of zero. Deviation to a higher price would also yield zero profit, due to the binding auction selection rule. Since $\rho \leq 1 - c(q^L)/v(q^H)$, withdrawing from the market would not improve their position. For all bids strictly above $c(q^L)/(1-\rho)$, a potential seller has an incentive to unilaterally undercut his competitors in the continuous bid interval: by doing so, he can increase the probability of trade from $1/n$ to one. The full trade surplus accrues to the buyer. **B.** For higher retainage proportions, $\rho > 1 - c(q^L)/v(q^H)$, there exists no price in the available bid interval that would induce a seller to participate. Thus, the only rationalizable action in the Bidding Stage is b^0. Participation would enable the buyer not only to appropriate the transaction

[10] Other Nash equilibria exist at higher prices and associated with non-participation, but since none of these outcomes are compatible with high quality, we ignore them here.
[11] For the special case of zero retainage, the Payment Stage is obsolete.

surplus, but to appropriate part of the seller's cost outlay. Off the equilibrium path, the rationale for payment, production and procurement decisions are the same as in part A. ■

The standard theory tells us that below a retainage threshold determined by the ratio of a seller's minimum production cost to the buyer's maximum valuation for the product, the market will collapse to an inefficient (low quality) outcome. Above this threshold, the burden of trust is too high for sellers to participate and the market unravels.

2.2 *A Model of Fair Payment Norms*

In a second approach, we adopt a simplified version of Bolton and Ockenfels' (2000) ERC theory: the α model. Suppose again that potential sellers compete anonymously in the Bidding Stage and that in the Procurement Stage, the buyer accepts any project she expects to be profitable. A buyer's agreement to purchase initiates a fundamental transformation (Williamson 1985). Now suppose that, after the fundamental transformation takes place, fair payment concerns become salient in the buyer-seller relationship. A fixed proportion α of buyers in the population are concerned about making a fair payment to sellers for quality provided, to the extent that discretion is permitted in the contract. Denote those buyers as "trustworthy". The remaining proportion $(1-\alpha)$ of buyers care only about maximizing their own payoff, as before. Denote these buyers as "self-interested". Whether a buyer is trustworthy or self-interested is private information, but the proportion α is common knowledge to all agents.

This rather spare model enables approximation of the population propensity to reciprocate using a single parameter α.[12] Bolton and Ockenfels (2000) assume that the fair reference point in a bilateral relationship is the 50-50 surplus split.[13] We extend their α model to incorporate the observation that not all supply chain relationships are created equal (Cui et al. 2007), by permitting the fair reference point to vary.

Trustworthy buyers' payoffs for terminal nodes following the Procurement Stage reflect a preference to minimize the difference between the winning seller's realized surplus share received, $\sigma_S = \pi_S / s(q^j)$, and a reference fair surplus share, $\gamma > 0$.[14] The seller's realized surplus share can be negative. The parameter γ is exogenous and represents a fair payment norm, which in practice will be context-specific. Note that while the fairness norm is independent of the quality

---

[12] Bolton and Ockenfels (2000, p 167) succinctly explain the rationale for such an approach: "much of what we need to know has to do with the thresholds at which behavior deviates from the 'more money is preferred to less' assumption".
[13] When assuming a 50-50 fair reference point, related models of fairness (e.g. inequity aversion à la Fehr and Schmidt 1999) yield qualitatively similar results for bilateral trade relationships.
[14] Production costs are included in surplus share comparisons, based on prior evidence from hold-up experiments that buyers consider sunk costs when making decisions on final surplus divisions (Ellingsen and Johannesson 2004, Hackett 1994).

level, the payment action that achieves a payoff outcome consistent with this norm does depend on the price and quality.[15] Thus in the Payment Stage, a trustworthy buyer sets *r* so as to minimize the absolute distance $|r\rho p + (1-\rho)p - w(q|\gamma)|$, where the reference fair total payment $w(q|\gamma) = \gamma v(q^j) + (1-\gamma)c(q^j)$ is a function of the chosen product quality, conditional on the prevailing fairness norm.

There are some notable special cases in this game of incomplete information. The standard theory is captured by $\alpha = 0$. Interestingly, if $\alpha = 1$, sellers know that all buyers are trustworthy and so there is always an incentive to place either a marginally or substantially lower bid in the available interval, associated with low quality. The presence of self-interested types in the population is necessary for high quality to emerge. When $\rho = 0$, the model collapses to an extensive-form game of complete information and Proposition 1A applies. Thus, in fixed-price contracts, high quality is never an implementable outcome regardless of the level of population reciprocity.

Let us now assume that there is some degree of flexibility in the contract, i.e., $\rho > 0$, and sellers have incomplete information about a buyer's payment preferences, i.e., $\alpha \in (0,1)$. The resulting model can be analyzed as an extensive-form game of incomplete information. We apply a perfect Bayesian equilibrium solution concept and restrict attention to pure strategy equilibria. This requires that beliefs correspond to the objective probabilities for all equilibrium actions.

In the Payment Stage, a self-interested buyer always returns zero retainage. Let σ_S' denote a seller *i*'s interim surplus share, i.e. if the retainage return will be zero. If $\sigma_S' \geq \gamma$, a trustworthy buyer will also return zero retainage. If $\sigma_S' < \gamma$, she will return the retainage fraction necessary to achieve a profit distribution as close as possible to the fair reference point. The best-response payment function for a trustworthy buyer is defined in (1) and produces an interior or corner solution depending on the price and quality level.

$$r(p, q_i) = \begin{cases} 0, & if\ \sigma_S' \geq \gamma \\ \min\left\{\frac{w(q_i|\gamma)}{\rho p} - \frac{(1-\rho)}{\rho}, 1\right\}, & if\ \sigma_S' < \gamma \end{cases} \tag{1}$$

where $w(q_i|\gamma)$ is non-decreasing in seller *i*'s quality choice.

Corollary 1: *The retainage return fraction in exchange for high quality is weakly higher than the retainage return fraction for low quality.*

[15] In other words, γ is not an injunctive norm by which high quality is perceived as morally right and low quality as morally wrong. Instead, it represents a community - or industry - norm as to what constitutes fair payment for works provided.

In the Production Stage, the seller anticipates the retainage payment decisions of buyers on average and chooses a quality level to maximize his expected monetary payoff as follows:

$$E[\pi_S] = (1-\rho)p + E(r;p,q_i)\rho p - c(q^j) \ ,$$

where a seller's expectation of the retainage return fraction given a quality function and bid is driven by the posterior belief of the proportion of trustworthy buyers in the population,

$$E(r;p,q_i) = \alpha r(p,q_i). \tag{2}$$

We obtain a second corollary.

Corollary 2: *The lowest bid that can implement high quality is* $\underline{b}^H = \frac{c(q^H)}{1-\rho+\rho E(r;p,q^H)}$.

Proof. The expression is obtained directly by setting the seller's expected payoff function to zero and solving for the bid, where $E(r;p,q^H)$ is defined in (2). The bid $\underline{b}^H$ is decreasing in α and γ, and increasing in ρ. ■

The seller will prefer to produce high over low quality when the expected payment premium more than compensates him for the increase in production cost. We require ex post incentive compatibility (EPIC):

$$EPIC\!: \ \alpha \geq \frac{c(q^H) - c(q^L)}{(r(p,q^H) - r(p,q^L))\rho p} \ , \tag{3}$$

The numerator is the marginal cost of producing high quality. The denominator is the expected marginal return of producing high quality, i.e., the difference in retainage returns from a trustworthy buyer in exchange for high over low quality. At its binding level, EPIC is concave upward for sufficiently high auction prices and has a global minimum at a price equal to $w(q^H|\gamma)$, which splits the surplus according to the fair payment norm exactly.[16]

Since trade is voluntary, any equilibrium seller bidding strategy must also satisfy ex post participation (EPPC):

$$EPPC\!: \ \alpha \geq \max\left\{\frac{c(q^L) - (1-\rho)p}{r(p,q^L)\rho p}, \frac{c(q^H) - (1-\rho)p}{r(p,q^H)\rho p}\right\} \tag{4}$$

[16] At contract prices below $w^\mu(q^H)$, EPIC is non-increasing in price as $(r(p,q^H) - r(p,q^L))$ becomes larger. At contract prices above $w^\mu(q^H)$, EPIC is non-decreasing in price as $(r(p,q^H) - r(p,q^L))$ becomes smaller.

which is defined for all $\sigma_S' < \gamma$ and is non-decreasing in the retainage proportion. At all interim surplus shares greater than σ_S', EPPC is satisfied. We obtain a third corollary.

Corollary 3: *With a larger discretionary payment component, a higher proportion of trustworthy buyers is required for the market not to unravel.*

Proof. First, note that regardless of α and ρ, a necessary condition for a bid to be profitable in expectation is that it exceeds the seller's production cost. For all such bids, the numerators in (4) represent the additional compensation that a seller requires to break even after the Production Stage. This requirement is strictly increasing in ρ. Second, recall that a trustworthy buyer prefers the seller to receive $w(q_i|\gamma)$, which is strictly above the requirement stated in the relevant numerator and is independent of ρ. A trustworthy buyer's preference is constrained by the available retainage money ρp. For any given bid and quality level, an increase in the retainage proportion will not change the relevant denominator in (4) because the trustworthy buyer's retainage return adjusts to exactly, or less than, offset the increase in ρp. This can be seen from inspection of (1). ■

We can now characterize trade conditions under which there exists a perfect Bayesian equilibrium associated with high product quality.

Proposition 2. *If $\rho > 0$, $\alpha \in (0,1)$ and $\underline{b}^H \in [c(q^L), w(q^H|\gamma)]$, then there exists a perfect Bayesian equilibrium in pure strategies characterized as follows: $b_i^* = \underline{b}^H$ for all i, $a^* = a^1$, $q_i^* = q^H$, trustworthy buyers return $r^*(p, q_i)$ and self-interested buyers return $r^* = 0$.*

Proof. Suppose α and ρ are such that a choice of high quality is EPPC and EPIC at some available bid $b \leq w(q^H|\gamma)$. In the Procurement Stage, a buyer will accept a price in this region independently of whether she is trustworthy or self-interested, because it guarantees her a non-negative profit. The procurement decision is therefore uninformative, and the seller does not update his prior belief about α. In the Bidding Stage, there is no incentive for a potential seller to deviate to a higher bid because it would yield zero selection probability at auction. To ensure that there is no incentive for a potential seller to submit any lower bid from the available interval, we require that $E[\pi_S(b - \varepsilon, q^j)] - E[\pi_S(b, q^H)] \leq 0 \; \forall \, j \in \{L, H\}$. This condition is satisfied for q^H when α is at a weakly decreasing point on the EPIC, which from its concavity is true at all bids below $w(q^H|\gamma)$. Since high quality is EPIC and the EPPC (by inspection) is strictly decreasing in price, it must also be true for q^L. It follows from the monotonicity of the payoff functions and from Corollary 2 that the lowest bid that satisfies these conditions is $\underline{b}^H$. Thus, if $\underline{b}^H \in [c(q^L), w(q^H|\gamma)]$, then a seller bidding strategy composed of this bid combined with high quality is an equilibrium strategy. ■

The intuition for Proposition 2 is captured in Figure 2. When $\alpha = 0$, the winning seller's expected payoff functions, conditional on quality, are linear in the bid. With a strict mixture of trustworthy and self-interested buyer types in the population, the seller's payoff functions remain monotonic but are now non-linear in the bid. The more optimistic the commonly held belief about the level of population trustworthiness α, the wider the range of available seller bidding strategies (combination of bid and quality level) which generate a non-negative profit in expectation. If there exists a feasible belief on α and retainage proportion ρ for which a combination of bid $b_i < w(q^H|\gamma)$ and high quality breaks even, then there can be no incentive for a seller to undercut to any lower bid so long as a combination of the same bid with low quality would generate a loss. The result breaks down as α goes to 1, because the knowledge of a guaranteed fair payment restores the seller's marginal incentive to undercut his competitor until reaching the minimum production cost.

To demonstrate that Proposition 2 is interesting, consider the following example. Suppose that $n = 2$ potential sellers compete to win a procurement contract under the following cost and valuation schedules: $c(q^L) = 0.30$, $c(q^H) = 0.40$, $v(q^L) = 0.35$ and $v(q^H) = 0.80$ (units in tens of thousands). By setting $c(q^H) > v(q^L)$ we ensure that, with zero retainage provision, the buyer must accept vulnerability due to uncertain seller production.[17] In Figure 3, we depict equilibrium outcomes at all (α, γ) pairs in the unit square, for two non-zero retainage arrangements: $\rho = 0.50$ and $\rho = 0.75$. We will test these retainage arrangements in the laboratory experiment. Blue circles in the figure indicate (α, γ) pairs for which low quality is the equilibrium outcome predicted by Proposition 2. Green squares indicate (α, γ) pairs for which high quality is the equilibrium outcome predicted by Proposition 2. The numbers inside the shapes correspond to the equilibrium bid amounts.[18]

The figure demonstrates that if high quality is to be implementable, then given a retainage level, sellers must be sufficiently – but not overly – optimistic in their expectation of interacting with a trustworthy buyer. In the implementable regions, the more optimistic the belief about population trustworthiness α and/or payment norm γ, the (weakly) lower the equilibrium bid. The empty region in the right-hand panel is associated with seller non-participation in the high retainage arrangement ($\rho = 0.75$), when beliefs about population reciprocity are sufficiently pessimistic to deter a potential seller from entering a bid.

[17] This is in-keeping with Özer and Zheng's (2017) definition of trust in the supply chain.
[18] Recall for $\rho = 0$, trade is predicted to be inefficient (low quality) in this model independently of α and γ.

Figure 2. Seller bidding incentives in relation to the level of population trustworthiness α.

(a) $\alpha = 0$.

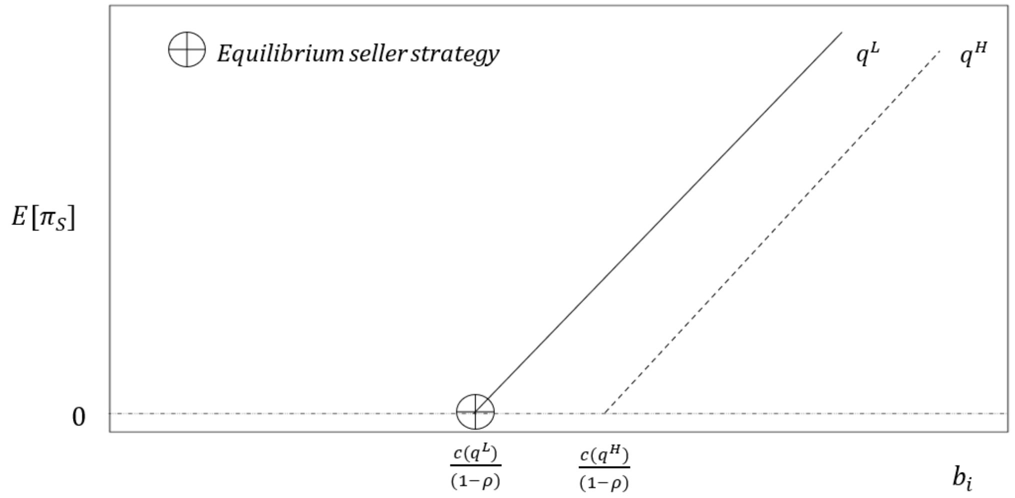

(b) $0 < \alpha < 1$.

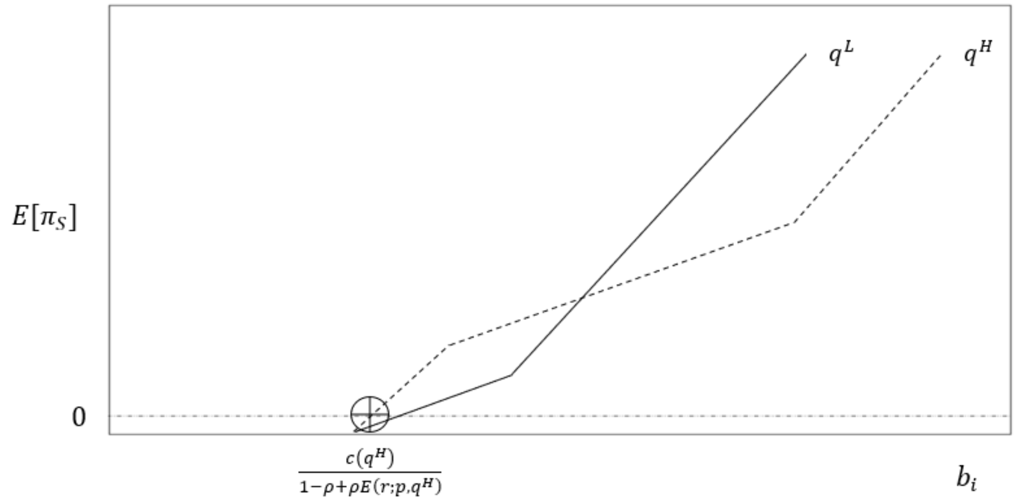

(c) $\alpha = 1$.

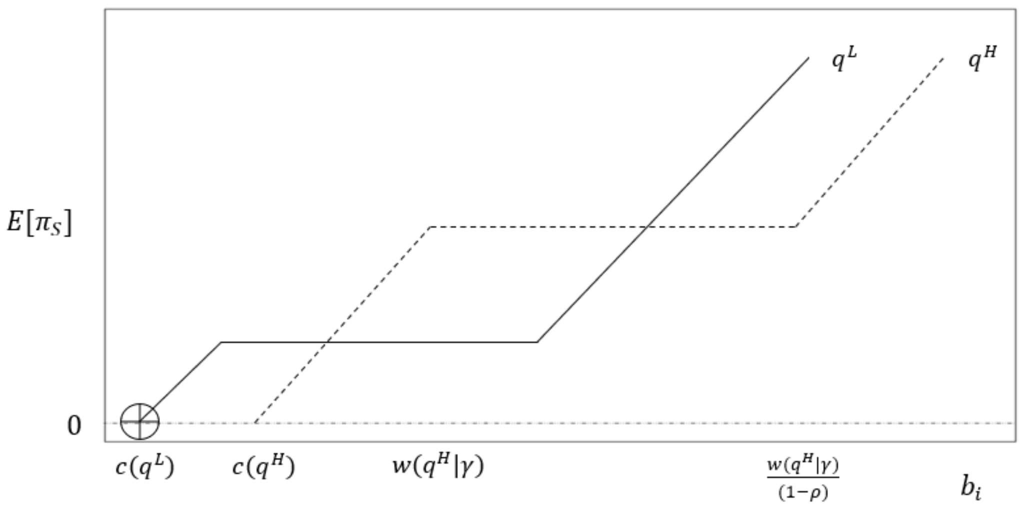

Figure 3. Parametric example: equilibrium seller strategies supported by the model of fair payment norms.

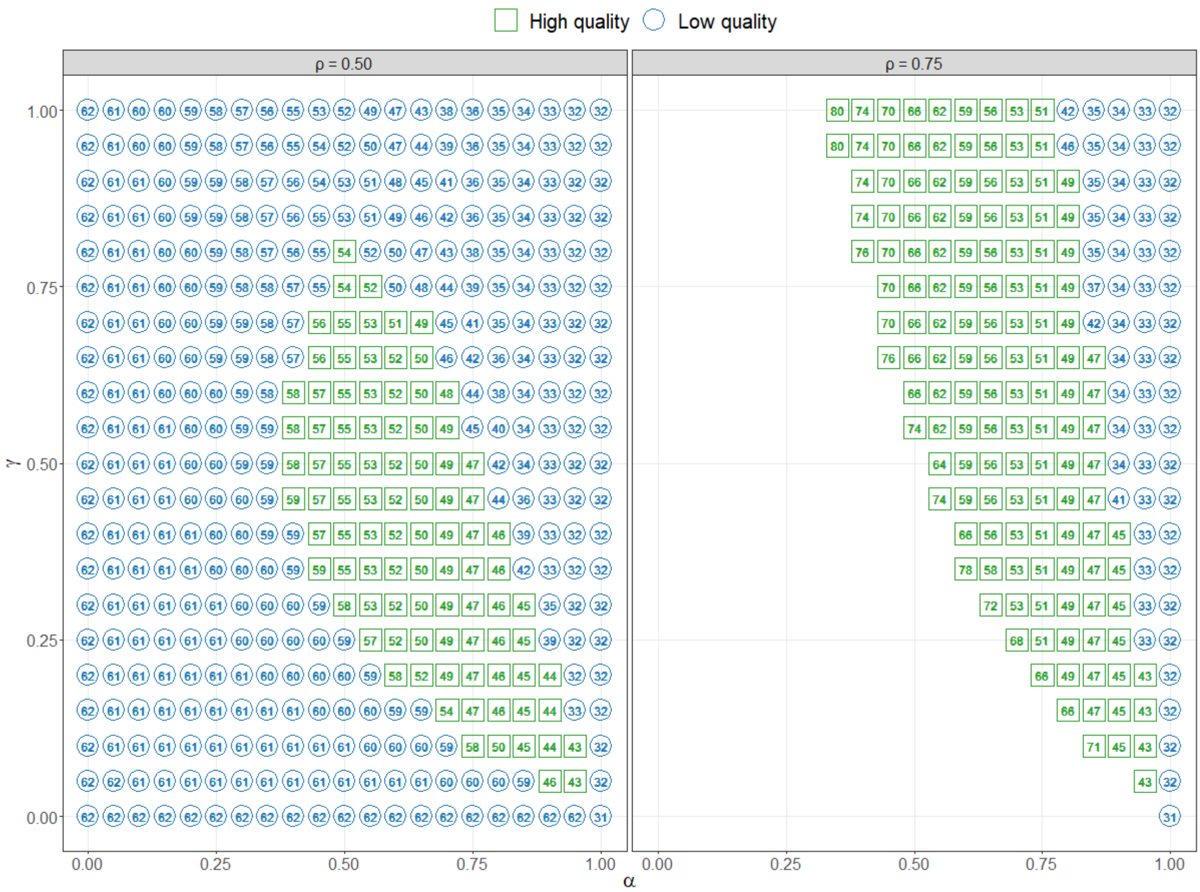

Notes: Figure displays the equilibrium seller bidding strategies as a function of population trustworthiness α and fair payment norm γ for the following levels of contractual retainage: $\rho = 0.50$ and $\rho = 0.75$. Based on the following cost and valuation schedules: $c(q^L) = 0.30$, $c(q^H) = 0.40$, $v(q^L) = 0.35$ and $v(q^H) = 0.80$. The numbers inside the shapes are the equilibrium bid amounts associated with the indicated quality level. An empty region corresponds to seller non-participation.

3. Experiment and Hypotheses

3.1 Experimental Design and Procedures

In our laboratory setting, we consider the case of two potential sellers. The valuation and cost parameters associated with high and low quality are indicated in Table 1. These values were displayed on the computer screens of all subjects. We restrict bids to be integers.

Table 1 – Experiment parameter values.

Quality Level $q \in Q$	$q = q^L$	$q = q^H$
Buyer Valuations: $v(q)$	35	80
Seller Costs: $c(q)$	30	40

We employed a between-subjects design with three treatments that varied the retainage provision, $\rho \in \{0, 0.50, 0.75\}$ – see Table 2. The $\rho = 0$ treatment benchmarks previous experiments in which payment of the full auction price is binding on the buyer, i.e., a fixed-price contract (cf. *Auction* treatment in Fugger et al. 2019). The $\rho = 0.50$ treatment simulates a setting in which potential sellers can fully offset the risk of partial or non-receipt of retainage monies by increasing their bids at auction. By contrast, in the $\rho = 0.75$ treatment, there is no bid available at which a potential seller can ensure to avoid a loss during trade. Thus, in this treatment, profitable trade can proceed based on the buyer's intent to compensate the seller for costs incurred – analogous to an LOI. We selected values $\rho = 0.50$ and $\rho = 0.75$ for the non-zero retainage treatments because they are the easiest retainage provisions for subjects to comprehend within the appropriate intervals from Proposition 1.

Table 2 – Experiment treatments.

Treatment	Retainage Level
1	$\rho = 0.00$
2	$\rho = 0.50$
3	$\rho = 0.75$

Subjects were randomly allocated to one of the three treatments and no subject participated in more than one session. Each treatment included six independent cohorts. There were three cohorts of the same treatment in every session. Each cohort consisted of three buyers and six sellers, who were randomly matched across 30 procurement interactions.[19] A bespoke algorithm guaranteed that no participant in a cohort played with the same pair of individuals in any two consecutive interactions. A total of 162 human subjects participated in our experiment sessions, which were conducted at the laboratory of a large public university in the United States. Participants were students recruited using web-based recruitment software.

All sessions followed the same protocol. Upon arrival, participants were seated at computer terminals and handed a written copy of the instructions to read in private. Terminals had physical dividers to prevent subjects from seeing the screens of other participants. The instructions were played from an audio recording at the front to ensure the description of the game was common knowledge and delivery consistent across sessions. The task was explained to subjects using a cover story related to the application of interest and the instructions included

[19] We did not inform subjects of the cohort size, to mitigate the possibility of tacit collusion in what might be considered a small cohort (see Katok 2011 for a discussion).

concrete terms such as "Buyer" and "Seller". This was a deliberate choice to improve subject understanding (see Cooper and Kagel 2003) and increase external validity (see Katok 2017).[20]

Participants completed a computerized test of understanding before being assigned to their role as a buyer or seller and matched into their first interaction group. Roles remained fixed throughout. Communication was prohibited and all interactions were anonymous. Own-group feedback was provided between periods. This information remained available in a history table to reinforce the game-theoretic assumption of "perfect recall". At the end of a session, participants answered a non-incentivized questionnaire to elicit demographic information, attitudes to trust and risk. The experimental interface was programmed using oTree software (Chen et al. 2016).

Subjects received monetary incentives for their participation. Each subject was paid his or her summed experiment earnings privately and in cash at the end of a session, in addition to a $5 show-up fee. Payment was made sequentially, with sufficient time intervals between any two subjects to mitigate against the possibility of side-payments. We used a symmetric exchange rate of 20 experiment currency units (ECU) to $1. Average subject earnings were $17.70 in the $\rho = 0$ treatment, $22.70 in the $\rho = 0.50$ treatment and $25.00 in the $\rho = 0.75$ treatment. Sessions lasted 60 to 75 minutes. Each subject received a non-refundable endowment of 7 ECU per period to cover potential losses. Subjects were informed that they would not leave the session with less than the show-up fee.[21] To reinforce the one-shot nature of interactions, we did not inform subjects about cumulative earnings until after the final period.

3.2 Hypotheses

Our theoretical predictions afford several testable experimental hypotheses, which are summarized in Table 3. Hypotheses 1 and 2 encapsulate the standard theory. In the $\rho = 0$ and $\rho = 0.50$ treatments, sellers will bid enough to cover their cost of producing low quality, competing away their profits in the process. With our experiment parameters, self-interest implies Nash transaction efficiency of 12.5%, measured as the percentage of surplus realized out of the total made available. The only difference hypothesized between treatments is the equilibrium price. Potential sellers are predicted to double their bids in the $\rho = 0.50$ treatment to compensate for anticipated loss of retainage money. In the $\rho = 0.75$ treatment, the retainage provision is too high for sellers to submit a profitable bid and so the market unravels.

[20] The instructions were framed in neutral language (see Zizzo 2010) and are contained in the Appendix B. We avoided the term "retainage". Audio recordings are available on request.

[21] In the $\rho = 0.75$ treatment, limited liability had to be imposed for two sellers. All results reported below hold if we exclude these two subjects from our analysis. In a pilot experiment, we tested the most extreme seller trust arrangement of 100% retainage. Seller losses, however, became a problem. Summary statistics for this variant are available on request.

We obtain alternative hypotheses from our model of fair payment norms. As shown in Figure 3, high quality is an equilibrium outcome in both the $\rho = 0.50$ and $\rho = 0.75$ treatments, conditional on sufficiently – but not overly – optimistic beliefs about the level of population trustworthiness α and the prevailing fairness norm γ. This confirms that Proposition 2 applies to our experiment environment. Hypotheses 3 and 4 follow directly from Corollaries 1 and 3: high quality will be rewarded reciprocally with higher retainage payment rates than low quality, and unraveling is most likely to occur in the $\rho = 0.75$ treatment. High quality is not implementable in the $\rho = 0$ treatment for any beliefs about α and γ. This is captured by Hypothesis 5.

Table 3 – Experimental hypotheses.

Hypothesis	Support
Standard Theory	
Hypothesis 1. For the $\rho = 0$ and $\rho = 0.50$ treatments: (a) Prices will not exceed $30/(1-\rho)$, which equals 30 and 60, respectively; (b) Transactions will fail due to buyer rejection at prices greater than $v(q^L)/(1-\rho)$, which equals 35 and 70, respectively; (c) Sellers will produce low product quality at minimum cost, independently of their bid; (d) Buyers will make zero retainage payment; and (e) Buyer profit will equal five and seller profit will equal zero.	Proposition 1A.
Hypothesis 2. For the $\rho = 0.75$ treatment, potential sellers will choose not to participate in the Bidding Stage and the market will unravel.	Proposition 1B.
Model of Fair Payment Norms	
Hypothesis 3. In the $\rho = 0.50$ and $\rho = 0.75$ treatments, retainage payment rates will be positively related to the product quality level.	Corollary 1.
Hypothesis 4. Market unraveling is most likely to occur in the $\rho = 0.75$ treatment.	Corollary 3.
Hypothesis 5. High quality is more likely to be observed in the $\rho = 0.50$ treatment and/or the $\rho = 0.75$ treatment than in the $\rho = 0$ treatment.	Proposition 2.

4. Experimental Results

4.1 Aggregate Findings

In Table 4, we present average values for the key outcome measures of our experiment.[22] Since there is no interaction between subjects playing in different cohorts, each cohort is considered a statistically independent observation. For each treatment, we have data from 540 procurement interactions.

As predicted by Hypothesis 1a, trade prices were significantly higher in the $\rho = 0.50$ treatment than in the $\rho = 0$ treatment (p-value = 0.001). While we statistically reject the point prediction that trade prices did not exceed 30 in the $\rho = 0$ treatment (p-value = 0.031), we fail to reject an average price of 60 in the $\rho = 0.50$ treatment (p-value = 0.563).

In contrast to Hypothesis 1b, there was a 31.5% buyer acceptance probability of prices above 35 in the $\rho = 0$ treatment, which were observed at 184 auctions. In the $\rho = 0.50$ treatment, a price above 70 was observed at just 14 auctions.

Conditional on buyer acceptance, Hypothesis 1c is supported in the $\rho = 0$ treatment. Low quality was chosen in 93.9% of transactions. There was a notable increase in market trust as represented by an improvement in average product quality in the $\rho = 0.50$ and $\rho = 0.75$ treatments. This cannot be explained by the standard theory.

We reject Hypothesis 1d in favor of our reciprocal Hypothesis 3. A seller's probability of receiving a non-zero retainage payment in the $\rho = 0.50$ treatment was 31% after choosing low quality and 61% after choosing high quality. Buyers paid to sellers 31.8% of retainage monies in exchange for high quality and just 6.1% in exchange for low quality. In the $\rho = 0.75$ treatment, where the retainage amount represents a larger share of the agreed price, the respective probabilities of non-zero payments were 54% and 74% and payment rates were 15.6% and 40.4%. The premiums paid for high quality are strongly significant (both p-values < 0.001).[23]

The standard theory does not provide a satisfactory explanation of the observed profit distributions. In qualitative terms, the predicted distribution of buyer-seller profits was reversed in the $\rho = 0$ treatment: sellers earned 1.60 more per period than buyers, a significant difference (p-value = 0.047). Thus, in this treatment, we reject Hypothesis 1e that sellers earn a profit of one (p-value = 0.031) and that buyers earn a profit of four (p-value = 0.063).

---

[22] We employ a two-tailed Wilcoxon Signed-Rank test for one-sample comparisons and two-tailed Wilcoxon-Mann-Whitney test for two-sample comparisons. When making multiple comparisons, we use Holm's (1979) p-value adjustment method. We acknowledge the potential *caveat* of arbitrary static correlations within sessions (Fréchette, 2012).

[23] We also find strong evidence for our reciprocal Hypothesis 3 in a regression analysis of individual-level decisions: high quality increases the average retainage payment rate on average by 24 and 30 percentage points in the $\rho = 0.50$ and $\rho = 0.75$ treatments respectively (both p-values < 0.001).

Table 4 – Cohort means and standard deviations.

	$\rho = 0$	$\rho = 0.50$	$\rho = 0.75$
Price	35.08	57.79	57.21
	(1.72)	(5.42)	(8.08)
Seller Participation	0.99	0.98	0.64
	(0.01)	(0.03)	(0.26)
Buyer Acceptance	0.60	0.91	0.98
	(0.20)	(0.08)	(0.01)
High Quality	0.06	0.30	0.50
	(0.02)	(0.20)	(0.19)
Retainage Payment	N/A	0.15	0.29
		(0.11)	(0.14)
For Low Quality		0.06	0.16
		(0.07)	(0.08)
For High Quality		0.32	0.40
		(0.16)	(0.16)
Market Unraveling	0.00	0.00	0.22
	(0.00)	(0.00)	(0.26)
Transaction Efficiency	0.11	0.36	0.56
	(0.03)	(0.17)	(0.17)
Buyer Profit	2.83	15.83	31.40
	(1.07)	(7.65)	(2.49)
Seller Profit	4.43	-0.25	-8.75
	(1.54)	(1.60)	(5.71)

Notes: Displayed are mean (SD) values for the key experiment outcomes based on 6 independent cohort observations per treatment. *Price*, *High Quality* and *Retainage Payment* are conditional on at least one seller submitting a bid at auction and buyer acceptance of this price. *Market Unraveling* is a measure of the proportion of auctions which failed to attract a single bidder. *Transaction Efficiency* is a measure of the surplus captured in the period divided by the surplus that sellers had made available. Profit is expressed in ECU per period conditional on trade and excludes the endowment.

In the $\rho = 0.50$ treatment, profit outcomes were rather different. While sellers earned a profit no different from zero on average (p-value = 0.844), buyers earned 15.83, an economically large and significant increase versus the theoretical benchmark (p-value = 0.031). Sellers fared significantly worse in the $\rho = 0.75$ treatment, incurring an average loss of 8.75 as opportunistic buyers appropriated surplus to double their profits to 31.40.

The experiment data also falsify Hypothesis 2: during 78% of auctions in the $\rho = 0.75$ treatment, at least one seller submitted a bid for the buyer to accept or reject. Yet although the seller non-participation rate was only 1 to 2% in the $\rho = 0$ and $\rho = 0.50$ treatments, it was 36% in the $\rho = 0.75$ treatment. Markets only unraveled in the $\rho = 0.75$ treatment. This result supports Hypothesis 4. There is no significant difference in the average price between the non-zero retainage treatments. Unsurprisingly, buyer acceptance rates were significantly higher in these treatments than in the zero retainage arrangement (p-value = 0.013), reflecting the reduced buyer vulnerability to loss. In the $\rho = 0.75$ treatment, nearly all auction prices were accepted.

The proportion of transactions associated with high quality was 30.2% in the $\rho = 0.50$ treatment and 50.4% in the $\rho = 0.75$ treatment, although the difference between the two treatments is only marginally significant (p-value = 0.093). This suggests that the burden of trust required to realize high quality was indeed lower with a non-zero retainage provision, as predicted by Hypothesis 5. To gain greater insight into the distribution of trade outcomes among the three treatments, in Figure 4 we present the relative frequencies of trade by price and quality level, conditional on at least one bid being submitted and on buyer acceptance. Prices are grouped in intervals of 10.

Figure 4. Relative frequencies of trade by price group and product quality.

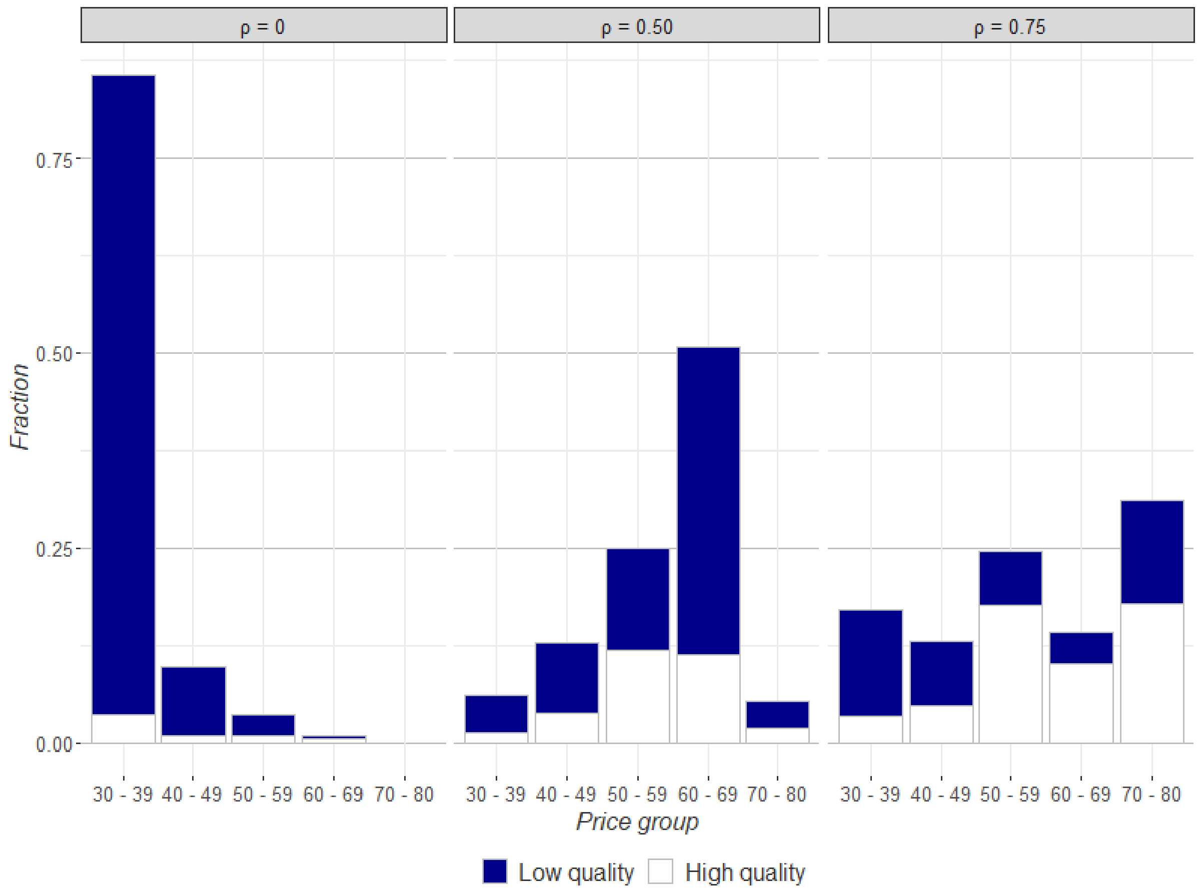

Notes: Based on 113 price-quality pairs across the experiment treatments. Relative frequencies are conditional on trade and do not reflect differences in the total number of transactions between treatments.

We observe in the figure that more than 85% of transactions in the $\rho = 0$ treatment were recorded in the price interval 30 to 39 and the quality of these transactions was near-uniformly low. More than 50% of transactions in the $\rho = 0.50$ treatment were observed in the price interval

60 to 69, and one in four of these was of high quality. Prices in this interval enabled the buyer to use her retainage return decision to impose the salient equal surplus split in exchange for high quality. A further 25% of transactions in the $\rho = 0.50$ treatment occurred in the interval 50 to 59 and around half of these were of high quality. In the $\rho = 0.75$ treatment, high quality was the majority choice for sellers at all prices over 50. That high quality was rarely observed at prices below 50 in the two non-zero retainage treatments is consistent with the range of implementable equilibrium seller bids shown in Figure 3.

Overall, we cannot reject the null hypothesis that on average, transaction efficiency attained its Nash equilibrium level of 12.5% in the $\rho = 0$ treatment (p-value = 0.31). There was, however, a substantial improvement relative to this level in the $\rho = 0.50$ treatment (36% transaction efficiency, p-value = 0.007) and in the $\rho = 0.75$ treatment (56% transaction efficiency, p-value = 0.009). This supports Hypothesis 5. In Figure 5, we plot a time series of the market unraveling complement and transaction efficiency across the 30 periods in our experiment. The trend differences in transaction efficiency are pronounced. In the $\rho = 0$ treatment, transaction efficiency fluctuated about its Nash equilibrium within the 0 to 25% interval. In the $\rho = 0.50$ treatment, transaction efficiency fluctuated within the 25 to 50% interval. The variability of this measure declined over time in both treatments. In the $\rho = 0.75$ treatment, transaction efficiency began in the 25 to 50% interval but then trended consistently upwards over time.

Market unraveling never occurred in the $\rho = 0$ and $\rho = 0.50$ treatments. In the $\rho = 0.75$ treatment, there was a marked unraveling in the proportion of auctions attracting at least one bidder after period five. This trend continued into the final period, at which point around half of markets unraveled. As a consequence, there is no significant difference in global efficiency – defined here as realized surplus out of total possible, rather than total made available by sellers – between the non-zero retainage treatments (p-value = 0.39). Notably, in the second half of the experiment, the difference in global efficiency between the $\rho = 0$ and $\rho = 0.75$ treatments is not significant at conventional threshold levels either (p-value = 0.13). By contrast, global efficiency remains significantly higher in the $\rho = 0.50$ treatment than in the $\rho = 0$ treatment after period 15 (p-value = 0.015).

4.2 Learning and Bidding Dynamics

Our theoretical analysis is equilibrium-based. Nevertheless, the aggregate results suggest scope for learning. We therefore conduct a formal analysis of the differences in decision-making and profit outcomes over time (see tables in the Appendix A).

Figure 5. Decomposition of market unraveling and transaction efficiency over time.

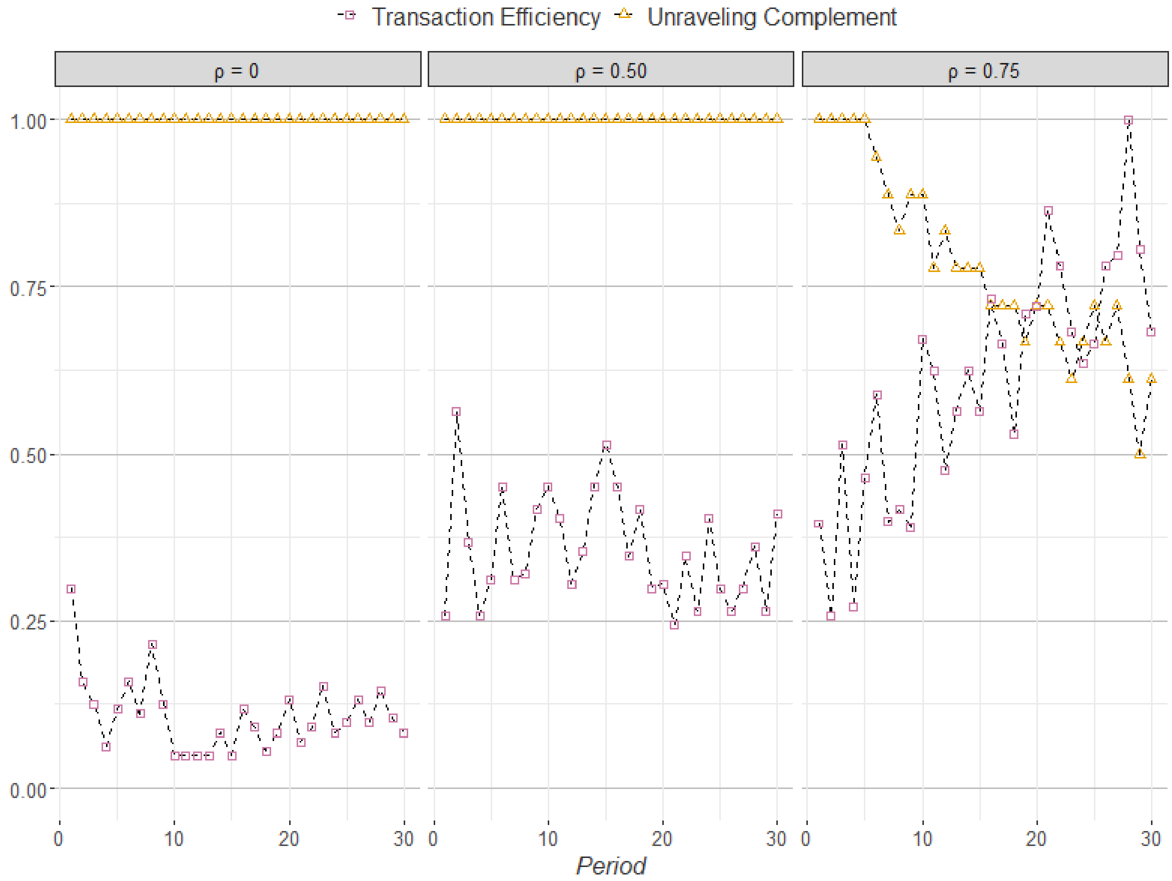

Notes: Based on 18 interaction groups per treatment in a period. *Unraveling Complement* is one minus the proportion of auctions which failed to attract a single bidder. *Transaction Efficiency* is a measure of surplus realized divided by surplus made available.

There are some statistically significant differences between periods 1 to 10 and 11 to 20. Sellers quickly learned to produce low quality in the $\rho = 0$ treatment. In the other two treatments, as buyers lowered their retainage payment rates, sellers learned to raise their bids or, exclusively in the $\rho = 0.75$ treatment, not to participate in the Bidding Stage. After the first 20 periods, there is no further significant bid adjustment towards the equilibrium prediction. In the $\rho = 0.75$ treatment, there is no evidence of a significant improvement in profits over time to warrant a choice of high quality by those sellers willing to participate. Any narrowing of the observed profit differential between buyers and sellers across period blocks is confined to the $\rho = 0$ and $\rho = 0.50$ treatments.

As a first step in analyzing sellers' bidding decisions, we code each seller non-participation decision as a bid at one increment above the highest available in the experimental price grid.[24] We then use a Tobit specification that censors the dependent variable from above

[24] It seems reasonable to suppose that non-participating sellers were unable to obtain a satisfactory guaranteed payment.

(Table 3). Standard errors are clustered at the subject level to account for the non-independence of observations.[25] We include the once lagged competitor's bid and once lagged trade as independent variables, along with a time trend and individual attitudes to risk and trust elicited in the post-experiment questionnaire.[26] Demographic information on subject age, gender, field of studies and income are used as control variables.

There is strong statistical evidence that bids exhibited a positive dependency in the most recently matched competitor's bid. In the $\rho = 0$ treatment, bids trended downwards over time. A reverse temporal bidding trend emerges in the $\rho = 0.50$ and $\rho = 0.75$ treatments, corroborating the aggregate learning analysis. In all three treatments, subjects who traded in the prior period submitted significantly lower bids on average in the following period. The smaller sample size in the third column reflects the elevated rate of seller non-participation in the $\rho = 0.75$ treatment, which results in the absence of lagged competitor's bids.

Table 5 – Censored regression analysis of seller bidding behavior.

	Bid		
	$\rho = 0$	$\rho = 0.50$	$\rho = 0.75$
Competitor's Bid $_{t-1}$	0.24***	0.51***	0.38***
	(0.05)	(0.08)	(0.06)
Trade $_{t-1}$	-4.68***	-6.35***	-5.76*
	(0.86)	(1.04)	(2.48)
Period	-0.16*	0.16***	0.55***
	(0.06)	(0.03)	(0.12)
Risk Index	0.08	0.60**	-1.58
	(0.43)	(0.20)	(1.15)
Trust Index	1.81	0.75	10.21*
	(2.80)	(1.79)	(4.59)
Constant	26.73***	29.98***	48.99***
	(5.90)	(8.76)	(12.78)
Control Variables	Yes	Yes	Yes
Log Likelihood	-3390	-3448	-2033
Observations	942	995	603
Right-censored	11	22	142

Notes: Standard errors in parentheses. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.
The models were estimated using a Tobit regression with robust standard errors clustered at the subject level. Seller non-participation decisions are coded as bids equal to 81 at which level the tobit specification right censors the dependent variable. Control variables: dummies for female and economics/business major, age and income rank (coefficient estimates for these control variables are available from the authors on request).

[25] The results are qualitatively unchanged if we instead cluster standard errors at the cohort level.

[26] Similar results are obtained for the $\rho = 0.50$ and $\rho = 0.75$ specifications if we include as a regressor the once-lagged retainage payment instead of the once-lagged trade indicator. Including both variables is not feasible due to multicollinearity. The risk and trust measures are non-incentivized and only suggestive. A few subjects did not provide a complete set of responses in the post-experiment questionnaire. The regression results are unchanged if we remove the control variables and use the full dataset.

4.3 *Representative Cohort Outcomes*

To obtain detailed insight into cohort heterogeneity, we plot procurement outcomes for each of the three interaction groups in every period by cohort. Figure 6 displays these outcomes for two representative cohorts per treatment, which suffice to capture the main behavioral differences that emerged in the experiment. The remaining four experiment cohorts per treatment exhibit a mixture of the representative cohort dynamics.

In Figure 6, the x-axis is the period-group and the y-axis is the winning auction price and/or total payment. An open triangle is an instance of market unraveling and a cross is an instance of transaction failure. A solid circle (square) is an accepted auction price at which the seller produced low (high) quality, and an open circle (square) is the corresponding total payment. The vertical arrows capture the differential between agreed price and total payment differential in the $\rho = 0.50$ and $\rho = 0.75$ treatments, i.e., when there is price flexibility in the contract. This approach yields 90 observations per cohort.

The top left-hand panel of the figure exemplifies the disciplining power of competition in the $\rho = 0$ treatment. After some initial learning, winning sellers in this cohort submitted bids in the 30 to 35 range and chose low quality. The top right-hand panel of the figure reveals a different dynamic: sellers attempted to elicit buyer trust and trade in the high-price region. Consistent with the standard theory, buyers rejected the majority of prices above 35.

The two middle panels in the figure demonstrate the convergence of winning bids in the $\rho = 0.50$ treatment towards the Nash equilibrium price of 60 from below. In cohort 5, sellers consistently produced low quality, receiving little if any retainage payment. In some other cohorts, sellers attempted to sustain reciprocal and high quality procurement. Buyers in cohort 3, for example, rewarded this choice with total compensation above the cost of high quality. This again supports our reciprocal Hypothesis 3. However, vulnerability to loss – as realized in interactions with a payment arrowhead below 40 – undermined market trust.

The bottom panels in the figure highlight opposing dynamics that emerged in the $\rho = 0.75$ treatment. Sellers in cohort 6 largely chose not to participate in the Bidding Stage, as predicted by the standard theory. Two-thirds of markets unraveled in this cohort. In contrast, sellers in cohort 1 were more willing to produce a high quality product. Where a buyer failed to reciprocate, the seller downside was substantially larger than in the $\rho = 0.50$ treatment (longer arrows in the figure), resulting in seller losses.

Figure 6. Representative cohort outcomes over time.

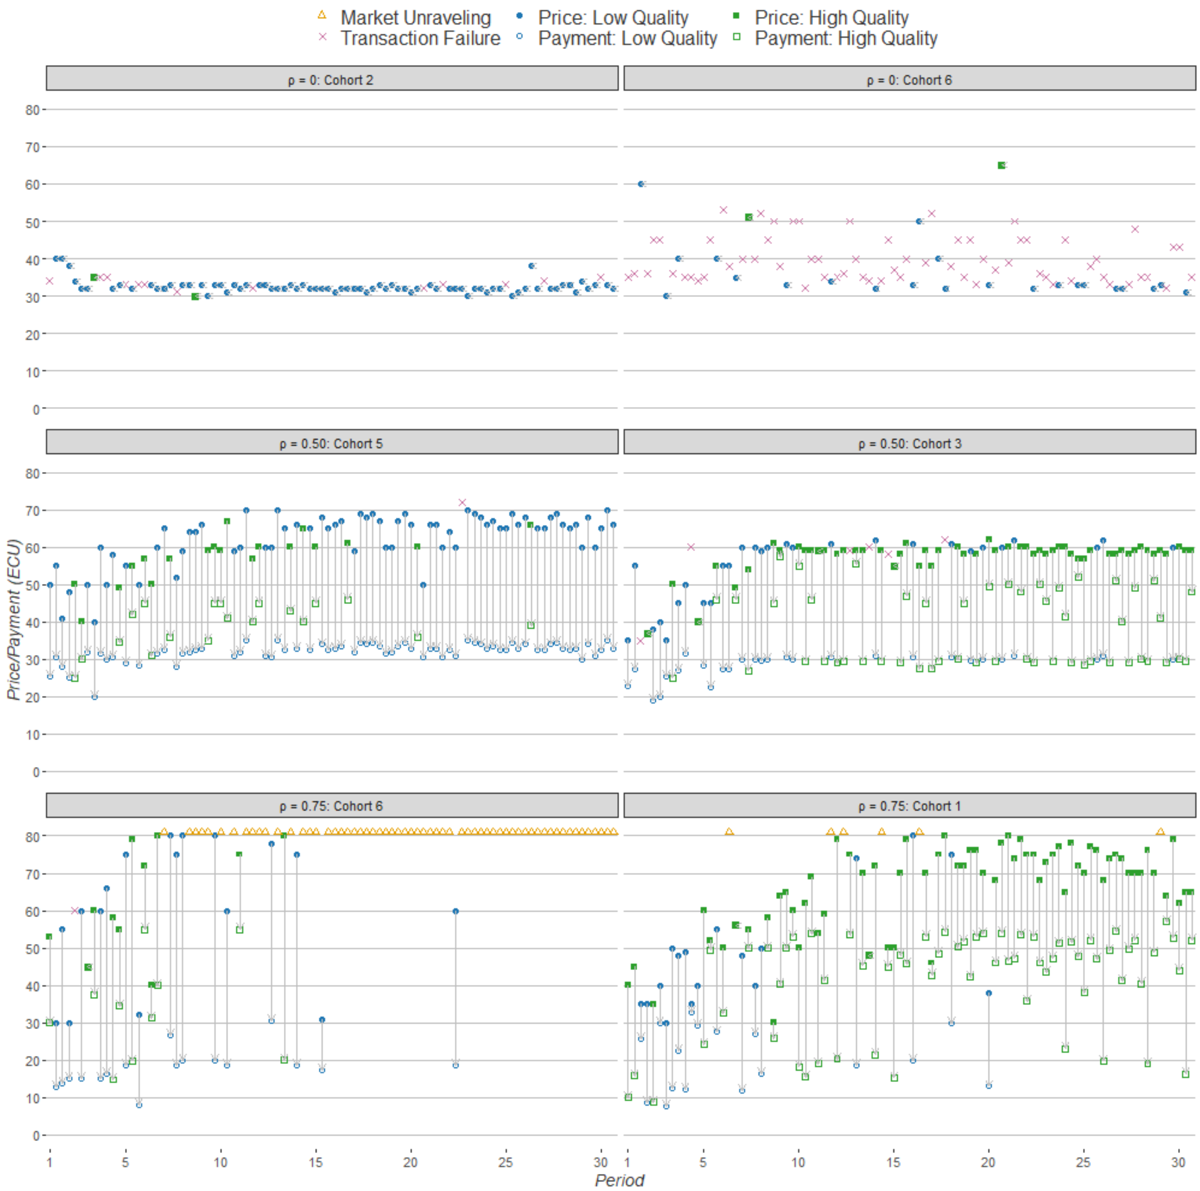

Notes: Panels display procurement outcomes over time in the specified treatment and cohort. Each observation corresponds to the outcome of a randomly matched interaction group in the period, with 1 buyer and 2 sellers. An open triangle is an instance of market unraveling, in which neither seller submitted a bid. A cross is a transaction failure, in which a buyer rejected the winning auction price. A solid circle is an accepted auction price at which the seller produced low quality. A solid square is an accepted auction price at which the seller produced high quality. An open circle and an open square are the corresponding total payments in instances where these differ from the auction price. In such instances, the vertical arrows represent the price-payment differential.

5. Estimating Trustworthiness and Fairness Norms

In our model of fair payment norms, we reduced the buyer population to two types: trustworthy and self-interested. This simplification enabled us to characterize outcomes along the dimensions of population trustworthiness α and fair payment norm γ, for different contractual arrangements. High product quality was rarely observed in the $\rho = 0$ treatment, consistent with the predictions of the model. In the $\rho = 0.50$ and $\rho = 0.75$ treatments, high quality was observed significantly more frequently, consistent with a range of equilibrium bidding strategies that vary with α and γ. While these parameters are latent, we can estimate the feasible set of joint parameters from buyer payment decisions in these treatments. To do this, we employ a tractable least absolute deviation (LAD) approach, which Bolton and Ockenfels (2000) used to apply their α model to behavior in gift-exchange and trust game experiments.

First, we define a procurement "bucket" $g_i(\rho, p, q)$ as a group of procurement interactions that share the same retainage provision and price-quality outcome. In terms of the strategic interaction, ρ determines the specification of the extensive form game, (p, q) specifies a penultimate information set and g_i identifies the subset of observations reaching this information set. Observations within this subset vary by the buyer's ultimate action.

For each bucket, we compute the average actual buyer and seller profits observed in the experiment.[27] We assign each bucket a weight according to its relative empirical frequency, i.e. the number of transactions in the bucket divided by the total number of transactions across all buckets (= 907). For each bucket, we then calculate the buyer and seller profits predicted by the model of fair payment norms at all (α, γ) pairs in the unit square. This enables us to compute, for each bucket and pair, the absolute deviation between actual and predicted buyer and seller profits. Finally, we take the weighted sum of these absolute deviations across all buckets and for each (α, γ) pair. The following objective function described this calculation:

$$\sum_i \frac{\#g_i}{907} \left(|\bar{\pi}_B^{actual}(g_i) - \bar{\pi}_B^{fair}(g_i, \alpha, \gamma)| + |\bar{\pi}_S^{actual}(g_i) - \bar{\pi}_S^{fair}(g_i, \alpha, \gamma)| \right) \quad (5)$$

Results of the estimation are presented in Figure 7. The subset of joint parameter estimates yielding the minimum LAD are identified by the darkest shaded color in the figure. This subset is centered around $\alpha = 0.40$ and $\gamma = 0.50$. Our best estimate is thus that 40% of buyers in the experiment were trustworthy, and those buyers perceived an equal share of the surplus as

[27] Taking averages reduces noise at the cost of losing information from higher moments. The results of the LAD estimation are qualitatively unchanged if we instead conduct the exercise at the transaction level. The least deviation estimate of trustworthiness is the same at $\alpha = 0.40$. The estimate of γ using the transaction level data is less precise, but also covers the best estimate found using the averaged data.

fair. Payoff equality is normatively appealing (Konow 2003), and equal surplus sharing is commonly found at the firm level (see Andreoni and Bernheim 2009).

We benchmark our estimates against previous studies. Using an analogous approach, Bolton and Ockenfels (2000) estimate a value of $\alpha = 0.50$ in the gift-exchange experiments of Fehr et al. (1993) and 0.42 in the trust game experiments of Berg et al. (1995). The best estimates of trustworthiness in our experiment then lie closest to trust game experiments. This is unsurprising; the post-auction phase of our procurement experiment has a trust-game flavor. Our estimate of γ predicts zero retainage payment for all interim seller surplus shares greater than one half. We observed this outcome in 137 out of 139 possible cases. Bolton and Ockenfels' ERC model specifies $\gamma = 0.50$ by assumption.

Figure 7. LAD estimates of the fairness model parameters.

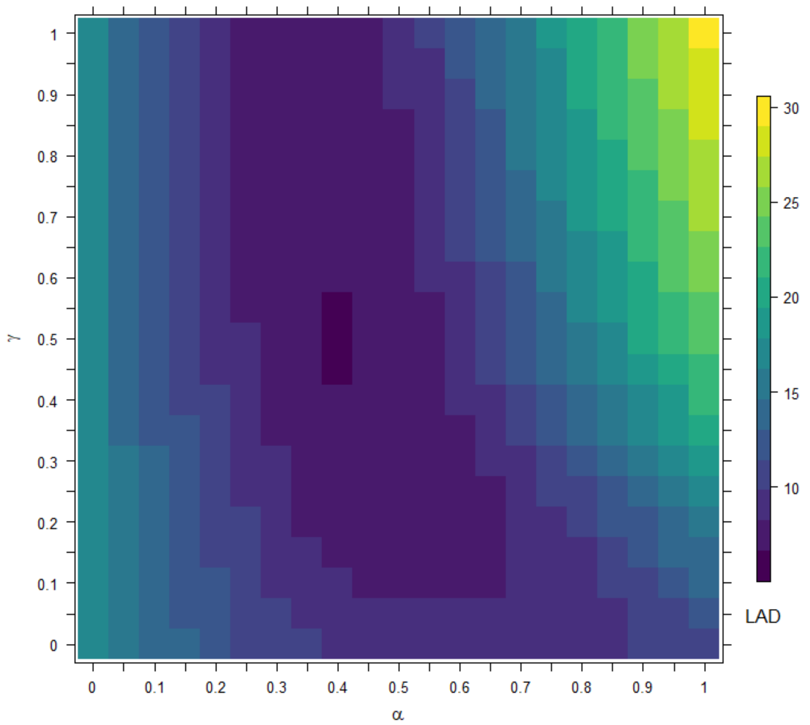

Notes: The heatmap presents the set of feasible joint parameter estimates of population trustworthiness α and fair payment norm γ, imputed from the experiment data. The color scale indicates the least absolute deviation (LAD) between average actual profits and those predicted by the model of fair payment norms (the unit is ECU). The darkest shaded region corresponds to joint parameter estimates that yield the minimum LAD between model's predictions and realized outcomes.

6. Concluding Remarks

The managerial literature has devoted substantial attention to how auction institution design can improve efficiency for procurement. In first price reverse auctions, standard principal-agent models show that high product quality equilibria do not exist due to moral hazard. The identification of alternative mechanisms that can reap the benefits of price competition while sustaining cooperation is of considerable managerial value. We develop a model of fair payment norms and show that flexibility in the contract price through retainage incentives can help to circumvent the inevitable degradation of seller quality, as long as prevailing payment norms induce the anticipation of reciprocal buyer payments.

In our experiments, we find that the probability of high product quality is significantly increased in non-zero retainage arrangements relative to the benchmark zero retainage case. We also observe a marked trade-off between transaction efficiency and transaction fairness, due to the difficulty in sustaining trustworthy buyer payment behavior over time. When sellers are engaged informally, with no payment guarantee to cover production costs, we observe a marked unraveling in markets over time. This reduces procurement efficiency, despite the higher quality of transactions remaining in the market.

The experiment results reinforce the finding of Fugger et al. (2019) that reverse auctions perform poorly from an efficiency perspective with rigid fixed-price contracts. The benefits of mechanisms that permit buyers to incorporate aspects other than price in their procurement decisions have been discussed in detail (see, e.g., Englebrecht-Wiggans et al. 2007). Here we demonstrate that attention should also be given to the flexibility of the contractual price, even when price cannot be ex ante conditioned upon product quality.

The main message for procurement managers and construction industry practitioners is as follows. Retainage mechanisms, and deferred payments more generally, have a useful role in incentivizing sellers to provide high quality. This recommendation comes with two *caveats*. First, informal procurement arrangements, such as LOIs, can end up having an anti-competitive effect by discouraging potential sellers from participating in the market. Second, if retainage monies are not properly administered, sellers are left vulnerable to the payment malpractices of untrustworthy buyers. This can undermine seller profitability relative to fixed-price contracts and the deductibility of retainage throughout the supply chain would likely amplify these negative effects. Practitioners should remain cognizant of industry payment norms when designing contractual incentives.

Our study suggests that future research might usefully examine the effects of incorporating dispute resolution into the procurement setup, for example via arbitration. In the model outlined here, sellers are endowed with no bargaining power after the auction and cannot dispute the buyer's retainage return decision. This assumption is clearly restrictive. In a consumer setting, Andreoni (2018) observes that some legal enforcement is required to induce buyer trust when goods are sold with a "satisfaction guarantee". Nevertheless, formal litigation is often costly and alternative dispute mechanisms under-developed. Consideration could also be given to the impact of a buyer's payment record in the procurement process. Reputation has been observed to improve seller trustworthiness in reverse auctions with moral hazard (Brosig-Koch and Heinrich 2014).

We also note that the retainage-type incentives examined in this study are a feature of performance bonuses in principal-agent relationships more generally. Employment contracts often contain a fixed base payment, with the promise of discretionary rewards in the future (see also Fehr et al. 2007, Lee et al. 2018). In certain industries, such as finance or professional sports, these voluntary bonus payments may be the most lucrative part of the compensation package. A similar mechanism is used in the rental housing market by landlords, who withhold deposits as a percentage of the total rental price – typically amounting to one or two month's rent – to incentivize tenants to take good care of their property.[28] When there is an excess supply of agents, the motivating effects of performance bonuses are not obvious (MacLeod and Malcomson 1998). We provide evidence from a competitive bidding environment to suggest that implicit incentives can help mitigate the moral hazard problem, under such conditions of market imbalance.

Finally, an open question is the extent to which models of fairness, originally developed to explain personal exchanges, apply to the more impersonal setting of firm decision-making. We find that deriving common thresholds as to the likelihood of a buyer or seller in the market acting in good faith is a useful one. We acknowledge, however, that in repeated interactions the reference point may be path-dependent. This possibility is not captured by our static equilibrium model and is a productive avenue for further theoretical work.

[28] Arrangements to pay "half now, half later" also seem to be popular with criminal and gangster organizations, at least as portrayed in the movies.

References

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

Anderhub, V., Engelmann, D., & Güth, W. (2002). An Experimental Study of the Repeated Trust Game with Incomplete Information. *Journal of Economic Behavior and Organization, 48*(2), 197-216.

Andreoni, J. (2018). Satisfaction Guaranteed: When Moral Hazard Meets Moral Preferences. *American Economic Journal: Microeconomics, 10*(4), 159-189.

Andreoni, J., & Bernheim, B. D. (2009). Social Image and the Fifty-Fifty Norm: A Theoretical and Experimental Analysis of Audience Effects. *Econometrica, 77*(5), 1607-1636.

ASA (2018). Retainage Laws in the 50 States 2018. Alexandria, VA: Foundation of the American Subcontractors Association, Inc.

Bajari, P., Houghton, S., & Tadelis, S. (2014). Bidding for incomplete contracts: An empirical analysis of adaptation costs. *American Economic Review, 104*(4), 1288-1319.

Bajari, P., & Tadelis, S. (2001). Incentives versus Transaction Costs: A Theory of Procurement Contracts. *RAND Journal of Economics, 32*(3), 387-407.

Bausman, D. C. (2004). *Retainage practice in the construction industry*. Alexandria, VA: Foundation of the American Subcontractors Association, Inc.

Beer, R., Ahn, H.-S., & Leider, S. (2018). Can trustworthiness in a supply chain be signaled? *Management Science, 64*(9), 3974-3994.

BEIS (2017). Retentions in the construction industry. UK Department for Business, Energy & Industrial Strategy, Research Paper 17: Prepared by Pye Tait Consulting. Retrieved from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/654399/Retention_Payments_Pye_Tait_report.pdf

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior, 10*(1), 122-142.

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review, 90*(1), 166-193.

Brosig-Koch, J., & Heinrich, T. (2014). Reputation and mechanism choice in procurement auctions: An experiment. *Production and Operations Management, 23*(2), 210-220.

Chang, W.-S., Salmon, T. C., & Saral, K. J. (2016). Procurement auctions with renegotiation and wealth constraints. *Economic Inquiry, 54*(3), 1684-1704.

Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance, 9*, 88-97.

Cooper, D. J., & Kagel, J. H. (2003). The impact of meaningful context on strategic play in signaling games. *Journal of Economic Behavior & Organization, 50*(3), 311-337.

Cox, J. C., Isacc, R. M., Cech, P., & Conn, D. (1996). Moral Hazard and Adverse Selection in Procurement Contracting. *Games and Economic Behavior, 17*(2), 147-176.

Cui, T. H., Raju, J. S., & Zhang, Z. J. (2007). Fairness and channel coordination. *Management Science, 53*(8), 1303-1314.

Davis, A. M., & Hyndman, K. (2018). An Experimental Investigation of Managing Quality Through Monetary and Relational Incentives. *Management Science, 64*(5), 2345-2365.

Davis, A. M., & Leider, S. (2018). Contracts and Capacity Investment in Supply Chains. *Manufacturing & Service Operations Management, 20*(3), 403-421.

Ellingsen, T., & Johannesson, M. (2004). Promises, threats and fairness. *The Economic Journal, 114*(495), 397-420.

Elmaghraby, W. (2007). Auctions within e-sourcing events. *Production and Operations Management, 16*(4), 409-422.

Emiliani, M., & Stec, D. (2005). Commentary on "Reverse auctions for relationship marketers" by Daly and Nath. *Industrial Marketing Management, 34*(2), 167-171.

Engelbrecht-Wiggans, R., Haruvy, E., & Katok, E. (2007). A Comparison of Buyer-determined and Price-based Multiattribute Mechanisms. *Marketing Science, 26*(5), 629-641.

European Commission (2009). International Accounting Standard 11: Construction Contracts. Retrieved from https://www.ucetni-portal.cz/stahnout/ias-11-en_842.pdf.

Fehr, E., & Falk, A. (1999). Wage rigidity in a competitive incomplete contract market. *Journal of political Economy, 107*(1), 106-134.

Fehr, E., Gachter, S., & Kirchsteiger, G. (1997). Reciprocity as a Contract Enforcement Device: Experimental Evidence. *Econometrica, 65*(4), 833-860.

Fehr, E., Kirchler, E., Weichbold, A., & Gächter, S. (1998). When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor Economics, 16*(2), 324-351.

Fehr, E., Kirchsteiger, G., & Riedl, A. (1993). Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics, 108*(2), 437-459.

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

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*(3), 817-868.

Fréchette, G. R. (2012). Session-effects in the laboratory. *Experimental Economics, 15*(3), 485-498.

Fugger, N., Katok, E., & Wambach, A. (2019). Trust in procurement interactions. *Management Science, 65*(11), 5110-5127.

Hackett, S. C. (1994). Is relational exchange possible in the absence of reputations and repeated contact. *Journal of Law, Economics and Organization, 10*(4), 360-389.

Hart, O. (1995). *Firms, contracts, and financial structure*. Oxford: Clarendon Press.

Hart, O., & Moore, J. (1988). Incomplete Contracts and Renegotiation. *Econometrica, 56*(4), 755-785.

Herweg, F., & Schmidt, K. M. (2017). Auctions versus Negotiations: The Effects of Inefficient Renegotiation. *RAND Journal of Economics, 48*(3), 647-672.

Herweg, F., & Schwarz, M. A. (2018). Optimal Cost Overruns: Procurement Auctions with Renegotiation. *International Economic Review, 59*(4), 1995-2021.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics, 6*(2), 65-70.

Hoppe, E., & Schmitz, P. (2011). Can contracts solve the hold-up problem? Experimental evidence. *Games and Economic Behavior, 73*(1), 186-199.

Hu, S., Wan, Z., Ye, Q., & Chi, W. (2017). Supplier behavior in capacity investment competition: An experimental study. *Production and Operations Management, 26*(2), 273-291.

Katok, E. (2011). Using laboratory experiments to build better operations management models. *Foundations and Trends in Technology, Information and Operations Management, 5*(1), 1-86.

Katok, E. (2017). Designing and Conducting Laboratory Experiments. In K. Donohue, E. Katok, & S. Leider (Eds.), *The Handbook of Behavioral Operations*. Hoboken, NJ: Wiley.

Katok, E., & Pavlov, V. (2013). Fairness in supply chain contracts: A laboratory study. *Journal of Operations Management, 31*(3), 129-137.

Konow, J. (2003). Which Is the Fairest One of All? A Positive Analysis of Justice Theories. *Journal of Economic Literature, 41*(4), 1188-1239.

Lee, Y. S., Ribbink, D., & Eckerd, S. (2018). Effectiveness of bonus and penalty incentive contracts in supply chain exchanges: Does national culture matter? *Journal of Operations Management, 62*, 59-74.

Loch, C. H., & Wu, Y. (2008). Social preferences and supply chain performance: An experimental study. *Management Science, 54*(11), 1835-1849.

MacLeod, W. B., & Malcomson, J. M. (1998). Motivation and Markets. *American Economic Review, 88*(3), 388-411.

Özer, Ö., & Zheng, Y. (2017). Trust and Trustworthiness. In K. Donohue, E. Katok, & S. Leider (Eds.), *The Handbook of Behavioral Operations*: Wiley.

Özer, Ö., Zheng, Y., & Chen, K.-Y. (2011). Trust in forecast information sharing. *Management Science, 57*(6), 1111-1137.

Raina, P., & Tookey, J. (2013). The perceptions of retentions as held by clients, contractors and subcontractors. Mimeo.

RTS Flexible Systems Limited v Molkerei Alois Müller GmbH (2010). 14 UKSC.

Shachat, J., & Tan, L. (2015). An experimental investigation of auctions and bargaining in procurement. *Management Science, 61*(5), 1036-1051.

Waehrer, K. (1995). A Model of Auction Contracts with Liquidated Damages. *Journal of economic theory, 67*(2), 531-555.

Wang, R. (2000). Bidding and Renegotiation in Procurement Auctions. *European Economic Review, 44*(8), 1577-1597.

Wevill, J. (2015). *Law in Practice: The RIBA Legal Handbook* (2nd ed.). London: RIBA.

Williamson, O. E. (1985). *The Economic Institutions of Capitalism*: Simon and Schuster.

Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics, 13*(1), 75-98.

Appendix A: Learning analysis (relates to Section 4.2 of the main text).

Table A. 1 – Cohort means and standard deviations of decision variables by period block.

	$\rho = 0$		
	Periods 1-10	Periods 11-20	Periods 21-30
Bid	42.49	40.69^{+}	38.24
	(4.64)	(4.69)	(4.20)
Seller Participation	0.98	0.99	0.99
	(0.03)	(0.01)	(0.01)
Buyer Acceptance	0.56	0.53^{+}	0.73
	(0.20)	(0.24)	(0.22)
High Quality	0.15*	0.03	0.02
	(0.07)	(0.05)	(0.04)
	$\rho = 0.50$		
Bid	57.61**	63.46^{+}	64.75
	(5.66)	(6.47)	(5.72)
Seller Participation	0.99	0.99	0.95
	(0.01)	(0.01)	(0.07)
Buyer Acceptance	0.94	0.90	0.89
	(0.03)	(0.08)	(0.15)
High Quality	0.31	0.35	0.24
	(0.08)	(0.23)	(0.33)
Retainage Payment for Low Quality	0.10*	0.05	0.03
	(0.10)	(0.08)	(0.05)
Retainage Payment for High Quality	0.40	0.27	0.33
	(0.25)	(0.17)	(0.09)
	$\rho = 0.75$		
Bid	55.22*	64.47	62.35
	(4.43)	(9.51)	(16.21)
Seller Participation	0.83*	0.59	0.50
	(0.16)	(0.31)	(0.37)
Buyer Acceptance	0.96^{+}	0.99	1.00
	(0.03)	(0.01)	(0.00)
High Quality	0.37	0.53	0.55
	(0.15)	(0.25)	(0.37)
Retainage Payment for Low Quality	0.19	0.12	0.09
	(0.09)	(0.07)	(0.09)
Retainage Payment for High Quality	0.39	0.36	0.35
	(0.18)	(0.24)	(0.20)

Notes: Displayed are mean (SD) values for the key decision variables based on 6 independent cohort averages per treatment. Matched pairs t-test (two-sided) comparing averages between adjacent blocks, $^{+}$ indicates significant difference between adjacent blocks at the 0.10 level * indicates significance at the 0.05 level; ** indicates significance at the 0.01 level.

Table A. 2 – Cohort means and standard deviations of Buyer and Seller profits by period block.

	$\rho = 0$		
	Periods 1-10	Periods 11-20	Periods 21-30
Buyer Profit	4.16	0.97+	2.93
	(4.31)	(1.99)	(1.37)
Seller Profit	5.95	5.17+	2.89
	(3.19)	(2.27)	(0.75)
	$\rho = 0.50$		
Buyer Profit	17.99	16.77+	12.06
	(2.09)	(9.58)	(13.13)
Seller Profit	-2.30	0.40+	1.41
	(2.14)	(2.45)	(2.20)
	$\rho = 0.75$		
Buyer Profit	29.38	30.11	33.16
	(4.40)	(4.69)	(8.31)
Seller Profit	-11.34	-6.72	-8.90
	(3.06)	(7.56)	(9.13)

Notes: Displayed are mean (SD) values for the buyer and seller transaction profits based on 6 independent cohort averages per treatment. Matched pairs t-test (two-sided) comparing averages between adjacent blocks, + indicates significant difference between adjacent blocks at the 0.10 level * indicates significance at the 0.05 level; ** indicates significance at the 0.01 level.

Appendix B: Experimental instructions.

This is an experiment in decision making. If you read these instructions carefully and make good decisions, you may earn a considerable amount of money. The amount of money you earn will depend on both your decisions and the decisions of other participants. The currency unit in this experiment is called Experimental Currency Unit, or ECU for short. At the end of the session, you will be asked to complete a short questionnaire. Upon completion, your total earnings from the experiment will be displayed on the screen, including your participation fee of $5, and be paid to you in private and in cash.

How you earn money

For today's session, one-third of you will be randomly assigned a Buyer role, and two-thirds of you assigned a Seller role. You will see your role at the start of the session and this role will not change for the duration of the session. This experiment will include **30 rounds**. In each round a Buyer is matched with two Sellers. The Buyer can purchase a product from one of the two Sellers. The value of the product to the Buyer, and the cost that the Seller incurs to provide the product, depends on the Seller choice to deliver a High or Low Quality product, as follows:

	Low Quality	High Quality
Buyer Value	35	80
Seller Cost	30	40

Each player will start each round with 7 ECU. Therefore, when there is a transaction between a Buyer and a Seller,

- The Buyer earns: 7 + Profit, where Profit = Buyer Value – Payment to the Seller
- The Seller earns: 7 + Profit, where Profit = Payment to the Seller – Seller Cost

How a potential Seller is determined

An auction is used to select which of the two Sellers will have the opportunity to transact with the Buyer. At the start of each round, each Seller privately submits a bid for a price at which they would be willing to deliver the product or elects not to participate in the current round's auction. The Seller submitting the lower bid wins the opportunity to potentially trade with the Buyer. In the event of a tie, the computer will choose the winner at random. Bids can be integers from 30 to 80.

Example:

Seller A bids 35 and Seller B bids 62 = A is winner.

Sellers A and B both bid 40 = A and B each have 50% probability of winning.

A Seller who does not win or elects not to participate in the auction earns 7 ECU for the round (Profit = 0). After the auction, the Buyer is presented the winning bid. The Buyer chooses whether or not to proceed with the transaction. If the Buyer does not proceed the round ends and all three parties earn 7 ECU for the round (Profit = 0).

How the payment to the Seller is determined

Participants in the $\rho = 0$ *treatment received the following instructions:*

If the Buyer chooses to proceed, the Winning Seller receives **payment** equal to his or her bid amount. Then the Winning Seller must decide to deliver either a Low or High Quality product.

Examples:

- Suppose the winning bid is 50 and the Seller delivers Low Quality. The round profits for the Buyer and Winning Seller are as follows:

  Buyer Profit = 35 – 50 = –15*

  Seller Profit = 50 – 30 = 20

- Now suppose the Seller delivers High Quality. The respective profits are:

  Buyer Profit = 80 – 50 = 30

  Seller Profit = 50 – 40 = 10

**Note that it is possible to lose money in a round. Make your decisions carefully.*

Participants in the $\rho = 0.50$ *treatment received the following instructions:*

If the Buyer chooses to proceed, the Winning Seller receives **50%** of his or her bid amount. This is called the **initial payment**. Then the Winning Seller must decide to deliver either a Low or High Quality product. The Buyer learns the quality level of the product and then chooses how much of the remaining 50% of the winning auction bid is paid to the Seller. This amount is called the **deferred payment**. The Buyer keeps any portion of the remaining 50% of the bid not paid to the seller.

Examples:

- Suppose the winning bid is 50, the Seller delivers Low Quality, and the Buyer sets the deferred payment at 6. The round profits for the Buyer and Winning Seller are as follows:

  Buyer Profit = 35 – 25 – 6 = 4

  Seller Profit = 25 – 30 + 6 = 1

- Now suppose the Seller delivers High Quality and the Buyer still sets the deferred payment at 6. The respective profits are:

  Buyer Profit = 80 – 25 – 6 = 49

  Seller Profit = 25 – 40 + 6 = –9*

- Now Suppose the winning bid is 70, the Seller delivers High Quality and the Buyer sets the deferred payment at 30. The respective profits are:

  Buyer Profit = 80 – 35 – 30 = 15

  Seller Profit = 35 – 40 + 30 = 25

  **Note that it is possible to lose money in a round. Make your decisions carefully.*

Participants in the $\rho = 0.75$ *treatment received the following instructions:*

If the Buyer chooses to proceed, the Winning Seller receives **25%** of his or her bid amount. This is called the **initial payment**. Then the Winning Seller must decide to deliver either a Low or High Quality product. The Buyer learns the quality level of the product and then chooses how much of the remaining 75% of the winning auction bid is paid to the Seller. This amount is called the **deferred payment**. The Buyer keeps any portion of the remaining 75% of the bid not paid to the seller.

Examples:

- Suppose the winning bid is 60, the Seller delivers Low Quality, and the Buyer sets the deferred payment at 16. The round profits for the Buyer and Winning Seller are as follows:

  Buyer Profit = 35 – 15 – 16 = 4

  Seller Profit = 15 – 30 + 16 = 1

- Now suppose the Seller delivers High Quality and the Buyer still sets the deferred payment at 16. The respective profits are:

  Buyer Profit = 80 – 15 – 16 = 49

  Seller Profit = 15 – 40 + 16 = –9*

- Now Suppose the winning bid is 72, the Seller delivers High Quality and the Buyer sets the deferred payment at 44. The respective profits are:

  Buyer Profit = 80 – 18 – 44 = 18

  Seller Profit = 18 – 40 + 44 = 22

**Note that it is possible to lose money in a round. Make your decisions carefully.*

How you will be paid

At the end of the session the earnings from all rounds of the session will be converted to US dollars at the rate of 20 ECU for $1. These earnings will be added to your $5 participation fee, displayed on your screen, and paid to you at the end of the session.

Example experimental interface screenshot for sellers:

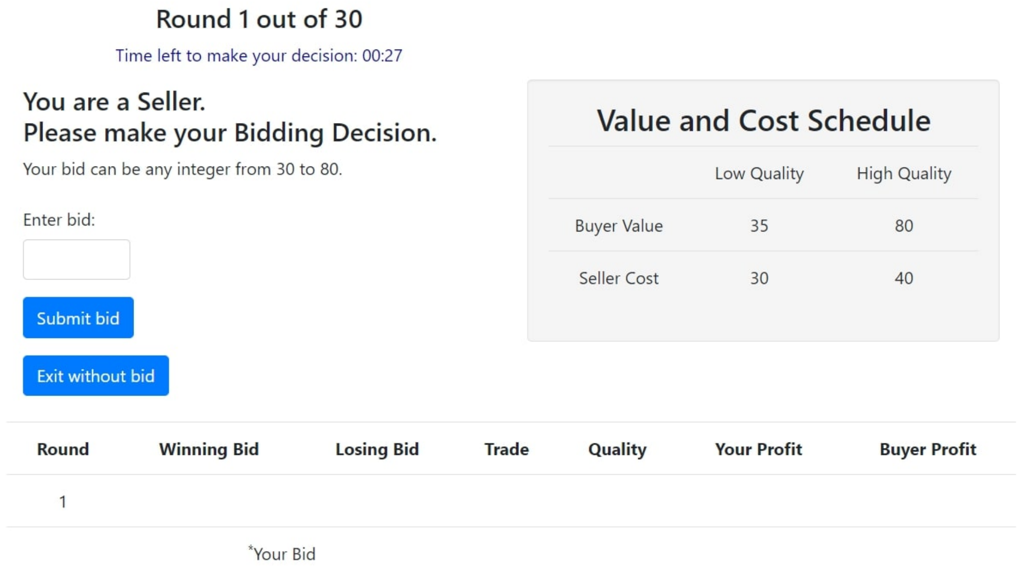

Example experimental interface screenshot for buyers:

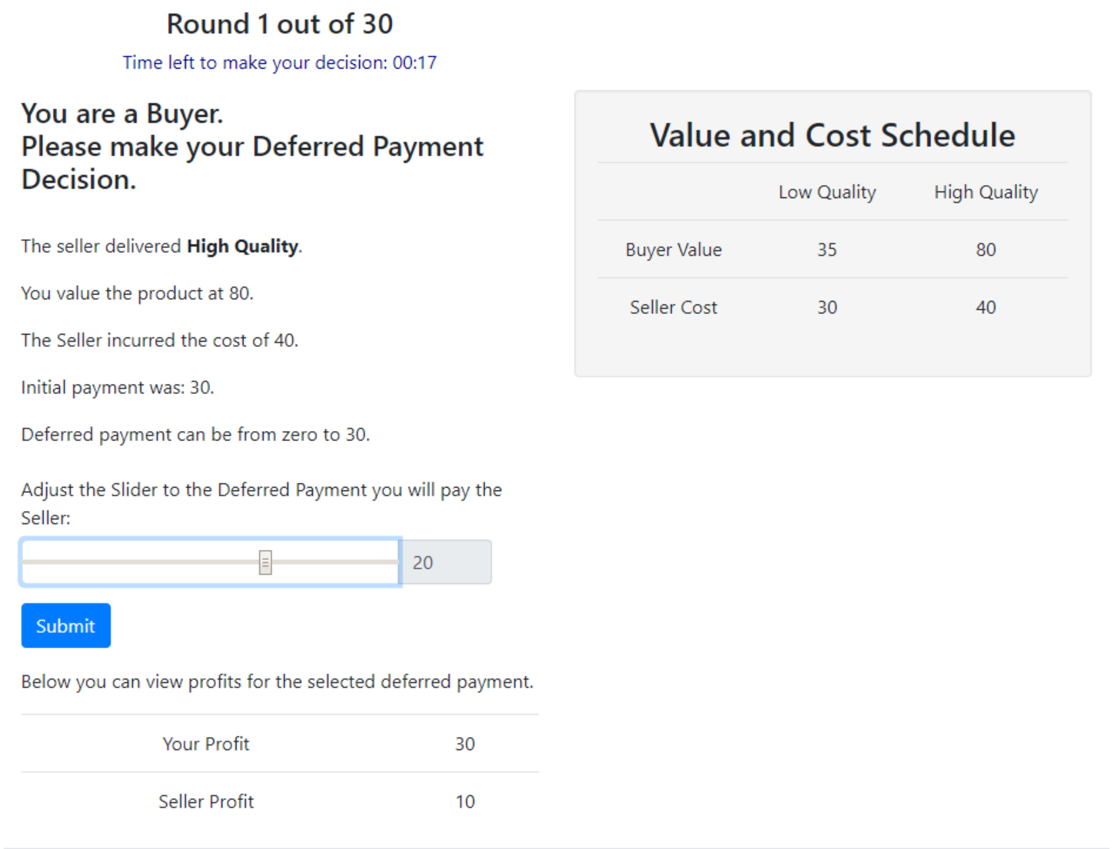

Note: In the $\rho = 0.50$ and $\rho = 0.75$ treatments, a slider with random initial value was used to avoid anchoring bias. A calculator displayed buyer and seller profits for the different available deferred payments.