KRANNERT GRADUATE SCHOOL OF MANAGEMENT

Purdue University
West Lafayette, Indiana

AN EXPERIMENTAL STUDY OF INFORMATION REVELATION POLICIES IN SEQUENTIAL AUCTIONS

by

Timothy N. Cason
Karthik N. Kannan
Ralph Siebert

Paper No. 1235
Date: June 7, 2010

Institute for Research in the
Behavioral, Economic, and
Management Sciences

An Experimental Study of Information Revelation Policies in Sequential Auctions

Timothy N. Cason* Karthik N. Kannan† Ralph Siebert‡

June 7, 2010

Abstract

Theoretical models of information asymmetry have identified a tradeoff between the desire to learn and the desire to prevent an opponent from learning private information. This paper reports a laboratory experiment that investigates if actual bidders account for this tradeoff, using a sequential procurement auction with private cost information and varying information revelation policies. Specifically, the Complete Information Policy, where all submitted bids are revealed between auctions, is compared against the Incomplete Information Policy, where only the winning bid is revealed. The experimental results are largely consistent with the theoretical predictions. For example, bidders pool with other types to prevent an opponent from learning significantly more often under a Complete Information Policy. Also as predicted, the procurer pays less when employing an Incomplete Information Policy only when the market is highly competitive. Bids are usually more aggressive than the risk neutral quantitative prediction, which is usually consistent with risk aversion.

JEL: C91, D44, D82.

Keywords: Complete and Incomplete Information Revelation Policies, Laboratory Study, Procurement Auction, Multistage Game.

Acknowledgments: We thank Dan Kovenock, Juan Carlos Escanciano, Sarah Rice, two anonymous referees and an associate editor, and seminar participants at the Universities of Calgary, Connecticut, Washington, and Melbourne, and Carnegie Mellon, Ohio State, Purdue and Southern Methodist Universities for helpful comments. Manish Gupte, Justin Krieg and Jingjing Zhang provided valuable research assistance. All errors are our own.

*Purdue University, Krannert School of Management, Department of Economics, 403 West State Street, West Lafayette, IN 47906-2056, USA, Email: cason@purdue.edu.

†Purdue University, Krannert School of Management, 403 West State Street, West Lafayette, IN 47907-2056, USA, Email: kkarthik@purdue.edu.

‡Purdue University, Krannert School of Management, Department of Economics, 403 West State Street, West Lafayette, IN 47906-2056, USA, Email: rsiebert@purdue.edu.

1 Introduction

In multistage non-cooperative games with information asymmetry, previous theoretical studies (e.g., Anand and Goyal, 2009, Kannan, 2010) have considered two learning aspects exhibited by the players – the incentive to extract and the incentive to obscure private information – as well as the trade-off between these incentives from the perspective of the market organizer. The managerial insights offered by these analyses depend on steep informational and rationality requirements for the players. Therefore, the relevance of these insights to practical, real-life situations remains an open question. In this paper, we investigate these learning aspects empirically using an experiment focused on one important application: procurement auctions.

Auctions are one of the most commonly-used mechanisms for procurement. General Dynamics, GE, Sears Logistics, and Staples, are among the many organizations that have employed auction technologies for procurement (Chandrashekar et al., 2007). These technologies enable a market-maker to easily alter the information policy for an auction; i.e., the extent to which information about bids are revealed at the beginning, during, and at the end of the auction. In fact, one of the important problems in the procurement context is the choice of the information policy (Elmaghraby and Keskinocak, 2000, highlights this issue). The problem has also been recognized in third party electronic procurement marketplaces (see, e.g., Jap, 2002, 2003, Arora et al., 2007). For example, in Freemarkets,[1] a firm that specializes in organizing electronic procurement markets, the buyer who convenes the market has a wide range of policy choices. At one end of the spectrum, the buyer can accept sealed bids and simply notify sellers individually whether each of them won or not. At the other end, all bids can be revealed as they are submitted, allowing bidders to respond in real time. Different information policies have also been adopted in traditional marketplaces. In federal and some state procurement auctions, the government is legally mandated to disclose only the winner's bid at the end of the auction (Milgrom and Weber, 1982). By contrast, in municipal construction contracting all bids are often revealed after the winner is selected (Thomas, 1996). Note that these procurement contexts typically feature repeated competition among the same suppliers across different auctions (e.g., Milgrom and Weber, 1982). Because of this repeated competition, the information revealed provides an opportunity for bidders to learn about their

[1] Freemarkets has merged with Ariba.

opponents' private information (such as their costs) across auctions. This opportunity can lead bidders to alter their behavior, which in turn affects the buyer's procurement costs.

In this context, we build on Kannan (2010) to model a procurement auction as a sequential private value auction in which winners do not drop out from subsequent auctions. Our model also permits bidders to have non-neutral risk preferences, since risk aversion has been documented as important in environments such as auctions, where agents face uncertainty related to the value of the object, the strategies used by others, and the private information possessed by their opponents. We use the model to study the following two policies in a first price sealed bid procurement auction: (i) the Incomplete Information Revelation Policy (IIP), in which only the winner's bid is revealed at the end of every auction, similar to the federal government mandated revelation policy; and (ii) the Complete Information Revelation Policy (CIP), in which all bids are made public at the end of every auction, similar to that in municipal construction auctions.

The perfect Bayesian Nash equilibrium analysis of the two policies identifies two key learning effects. The extraction effect, which occurs under IIP, refers to the bidders' desire to alter their bids so as to learn about the opponents' private information. The deception effect, which occurs only under CIP, refers to the bidders' desire to prevent their opponents from learning about their own private information by pooling with higher cost types. Although both these effects arise from a bidder's desire to maintain a relative informational advantage over her competitors, their influence on market outcomes can be different. Both effects lead to higher prices, but either may have a more dominant impact on procurement costs, depending on the degree of competition. Our analysis also presents insights into how the learning effects interact with risk aversion. While the bids in IIP decrease as the degree of risk aversion increases – a result consistent with previous findings, e.g., Maskin and Riley (1984) and Maskin and Riley (1987) – the analysis of CIP reveals a surprising result. We find that the pooling rate (submitting a high bid) actually increases with risk aversion.

The steep informational and rationality requirements imposed to compute the two stage Bayesian Nash equilibria are demanding, so the practical and descriptive value of the theory may be limited. This highlights the need to test the predictive power of the theoretical results. An obvious way to test the model would be through an empirical analysis of field data. However, private information about costs is typically not available, making such an analysis difficult. The choice of the information policy is also endogenous in the field, which complicates causal inferences.

An experimental study, through its use of a controlled setting and exogenous manipulation of information policies, can overcome these problems. Our experiment allows us to study bidding behaviors under different, exogenously-imposed information policies. The experiment focuses on the learning-related bidder behaviors which a buyer should consider when choosing procurement auction information policies in environments characterized by different degrees of competition.

The analysis of our experimental data shows that bidders pool with high-cost types to prevent opponents from learning significantly more often under CIP, and the procurer pays less under IIP only when the degree of competition is high. These results indicate that subjects behave as if they can compute a perfect Bayesian-Nash equilibrium of this dynamic game, or at least appreciate the intuition of the learning effects (the extraction and deception effects). We also observe that, in general, learning effects shift bids in the predicted direction with regard to treatment variations in the information policies and the degree of competitiveness in the market.

We also use the experimental data to apply a structural estimation procedure to estimate the degree of risk aversion most consistent with the bidding behavior. Our estimation results show that risk aversion is consistent with most of the overly-agressive bidding relative to the risk neutral Bayesian-Nash equilibrium observed. Bidders in the CIP treatment, however, do not pool with high-cost types at the high rates predicted by risk aversion. This may occur because cognitive limitations prevent subjects from understanding the risk reduction benefits of pooling, or because of non-pecuniary benefits of winning an auction. Thus, policy comparisons should also take these cognitive limitations and winning incentives into account.

Our paper is distinct in the literature in several ways. An extensive literature in auctions has analyzed the problem of information revelation in various contexts (e.g., Goeree, 2003, Das Varma, 2003, Katzman and Rhodes-Krop, 2008, de Silva et al., 2008, Chen and Vulcano, 2009). To the best of our knowledge, our paper is the first to experimentally study information revelation policies in a repeated procurement auction setting. Moreover, the notion of learning we focus on is across multiple stages within a game. It is different from prior experimental research which has extensively analyzed learning across multiple iterations of the same game (e.g., Camerer and Ho, 1999). Furthermore, the paper features theoretical, experimental, and structural estimation components. Thus, by drawing upon a diverse set of research methodologies, we provide a comprehensive analysis of the information policy choice problem. The analysis offers actionable managerial insights into the interaction between information disclosure policies and bidder learning effects. For example, there is a perception in the industry that bid transparency leads to aggressive bidding behavior and, therefore, lower buyer procurement costs, see e.g. Elmaghraby and Keskinocak (2000). However, the perception does not take into account obfuscation strategies which can be optimal with complete revelation of bids.

The rest of the paper is organized as follows. Section 2 reviews the related literature, and Section 3 summarizes the theoretical model and results. Section 4 presents the research hypotheses and the experiment designed to test them. Section 5 reports the basic results and main hypothesis tests, and Section 6 presents the structural estimation of the risk aversion model. Section 7 concludes.

2 Literature Review

2.1 Sequential Auctions with the Winner not Dropping Out

Suppliers in procurement auctions frequently compete against the same set of opponents across different markets or in different auctions held sequentially by the same buyer. Therefore, we model the procurement context as a sequential auction with the winner not dropping out. More specifically, the procurement setting is modeled as a private value auction, consistent with prior works (e.g., Maskin and Riley, 2000, Bajari, 2001).

Most prior work in the area of sequential auctions has focused on winners dropping out (e.g., Krishna, 2009, Klemperer, 2004). One of the first papers on sequential auctions with no winner drop out is Ortega-Reichert (1968), where a CIP-like policy is considered but in a common-value setting. A few other related papers have studied private value auction settings but have focused on results when the objects across the stages are stochastically equivalent (Branco, 1997, Englebrecht-Wiggans, 1994); have diminishing marginal valuations (Donald et al., 2006); or exhibit synergies (Jeitschko and Wolfstetter, 1998).

Hausch (1986) extends the Ortega-Reichert (1968) common value framework to compare the CIP-like sequential auction and the simultaneous auction. Two other closely related papers dealing with private value auctions are Thomas (1996) and Tu (2005). Thomas (1996) focuses on studying mergers, although part of his analysis compares policies that are similar to CIP and IIP. The model is similar to the case in which two bidders repeatedly compete and each is equally likely to be a low- or a high-cost type and their own cost type information is private. He concludes that CIP always generates higher buyer surplus and lower bids than IIP. This result is different from the conclusion in Kannan (2010), which shows that the surplus ranking of IIP or CIP depends on the distribution of cost types. In a recent work simultaneous to Kannan (2010), Thomas (2010) considers a distribution of cost types but continues to recommend CIP based on the analysis executed at the probability when the types are equally likely. Tu (2005) also analytically studies a two bidder game where bidders have private information about their cost types, except that their cost is drawn from a continuous distribution. He explicitly imposes an assumption that suppliers cannot pool with other cost types and determines that CIP generates higher buyer surplus than IIP. However, pooling is an important strategy we focus on in the present analysis.

2.2 Auction Experiments

Auctions have been studied extensively using experiments, including very practical applications such as for the design of FCC spectrum auctions (e.g., Goeree et al., 2006). Kagel and Levin (2010) provides a recent survey. Few experiments have focused on comparing the outcomes of different information revelation policies in sequential auctions, and only two studies focused on contexts similar to ours. Dufwenberg and Gneezy (2002) analyze the importance of information disclosure policy in a common value setup like Hausch (1986). The main difference, however, is that while Hausch (1986) considers first-price sealed-bid auctions, Dufwenberg and Gneezy (2002) consider a setup where bidders agree to share the royalty with the buyer. They consider three types of information revelation policies, including both of our policies where all bids, or all winning bids, are announced by the auctioneer between stages. They consider a common value setting, however, and bidders have no private value that could be potentially revealed from earlier bids, and competing bidders do not interact repeatedly. Theory does not predict any difference in bid prices between the bid revelation policies, but they find that when bidders are informed about the losing bids in previous stages, prices are significantly higher than the theoretical prediction. Bidders become more competitive when this information is not revealed, which moves bids closer to the theoretical prediction.

3 Theoretical Model, Equilibrium and Insights: Summary

In this section, we develop the theoretical results using a model similar to Kannan (2010), which considers a two-stage, private-value, sequential auction model with no winner drop outs. While

Kannan (2010) permits an arbitrary number of potential competitors, we examine the special case of two bidders (or suppliers) in our experiment. However, in the present model, bidders can be risk averse, while in the previous one they are all treated as risk neutral.

3.1 Model

Bidders are assumed to be characterized by a constant relative risk averse (CRRA) utility function of the form $U(x) = x^{(1-r)}$, where x is the payoff and r is the risk aversion coefficient. Thus, the risk neutral analysis is a special case of our model when $r = 0$. The CRRA approach "is the most widely used parametric family for fitting utility functions to data" (Wakker, 2008, p. 1329). We also assume that bidders have private costs which are drawn from a discrete distribution of two cost types.[2] Let c_l be the marginal cost of production for a low-cost supplier and c_h for a high-cost supplier, and θ be the probability with which a bidder is a low-cost type. We assume common knowledge of all three variables, although the outcome of the cost draw for each supplier is private information. Since we focus on bidders' learning across auctions, we model a two stage game, with each stage corresponding to an auction initiated by a buyer. The cost type for a bidder is determined before the beginning of the first stage and remains the same for both stages. In each stage, both suppliers simultaneously submit a sealed bid, $p \in [0, \infty)$. Only one winner is picked in each stage and he is the bidder with the lowest bid price in that stage (ties are broken randomly). The winner's payoff is his bid price minus his marginal cost, and the loser's payoff is zero. Therefore, the auction in each stage is a first-price sealed bid type. Between the first and the second stages, information is revealed according to the policy. Under CIP, all bids are revealed at the end of the first stage; while in IIP, only the winner's bid is revealed. We are interested in comparing the impact of the different information policies on bid prices and procurer surplus.

3.2 Equilibrium

The model specified above does not always have an equilibrium. Even single stage games, similar to the ones we encounter in the second stage, may not have an equilibrium. We follow Maskin

[2] Prior work has also used a two-type framework to study information policies in other settings, in part because some interesting aspects of learning across auction rounds do not exist in the continuous cost distribution case (Jeitschko, 1998). The main advantage of the two-type model is that the second stage is a relatively straightforward game and the analysis can focus on the learning effects in the first stage. This advantage is lost even in a three-cost type model, where the nature of the second stage equilibrium varies significantly depending on the first stage outcome, which must also be considered by bidders in the first stage.

and Riley (1985) to overcome the non-existence problem by implicitly assuming discrete bids in infinitesimal increments.[3] The game then has a unique Bayesian Nash equilibrium under CIP and IIP for each stage. We focus on the bidding behavior of the low-cost types since the bidding behavior of the high-cost type is uninteresting. A high-cost type always bids c_h in equilibrium independent of the policy. In the rest of this section, we summarize the equilibrium results.

3.2.1 Second Stage

Since the equilibrium is derived by backward induction, consider first the second stage. The second stage across the two policies has commonality, which can be captured as follows. Suppose one bidder, say A, believes that his opponent B is low-cost type with a probability of α, while B believes that A is low-cost type with a probability of β. Let $\alpha \leq \beta$, and suppose further that A already won the previous period with a price of p. The profit expressions for A and B when they are low-cost types and bidding a price of q are: $\pi_A(q) = (1 - \alpha F_B(q))\left((p - c_l + q - c_l)^{1-r} - (p - c_l)^{1-r}\right) + (p - c_l)^{1-r}$ and $\pi_B(q) = (1 - \beta F_A(q))(q - c_l)^{1-r}$, where $F_A(q)$ and $F_B(q)$ are the cumulative density functions (cdfs) of the bid distributions from players A and B. The equilibrium bid distributions are computed in Appendix A.1.

3.2.2 CIP Game

The nature of the equilibrium in the first stage game varies depending on θ and r. We first consider a separating equilibrium in the first stage.

Separating Equilibrium

Suppose a separating equilibrium exists in the first stage. This implies that the first stage reveals the type of the bidder. As a result, the second stage simply corresponds to a Bertrand game. The payoff across both stages from bidding p in the first stage is $\pi^{\text{CIP-sep},\theta}(p) = (1 - \theta)(c_h - c_l + p - c_l)^{1-r} + \theta(1 - F^{\text{CIP-sep},\theta}(p))(p - c_l)^{1-r}$, where $F^{\text{CIP-sep},\theta}(p)$ is the cdf. The equilibrium bid distribution is computed in Appendix A.2.1. Such an equilibrium is valid in the first stage under CIP only when $\theta < \frac{2^{1-r}-1}{2^{1-r}+1}$. When θ is larger than this threshold, the appendix also shows that only a semipooling equilibrium exists. Since $\frac{\partial}{\partial r}(\frac{2^{1-r}-1}{2^{1-r}+1}) < 0$, the threshold degree of competition (i.e., θ), above which

[3] Maskin and Riley's footnote 2 is directly applicable to our context: "As our model is formulated, an equilibrium in the sealed-bid auction may not exist. The nonexistence problem, however, is an artifact of our allowing literally a continuum of possible bids. In fact, we can restore existence even with a continuum by allowing the possibility of positive but infinitesimal bids, which we implicitly assume in our analysis."

the semipooling equilibrium exists, decreases with r.

Semipooling Equilibrium

A semipooling equilibrium in the first stage is one where the low-cost type pools with the high-cost type and bids c_h with a probability $\gamma < 1$, or submits a bid $p < c_h$, which reveals the low-cost type, according to the cdf $F^{\text{CIP-semi},\theta}(p)$. Let the corresponding pdf be $f^{\text{CIP-semi},\theta}(p)$ and the infimum of the distribution be p_l^{CIP}.

In the second stage game, three different possibilities exist: (a) When both bidders reveal their type to be low-cost, the second stage game is a Bertrand game. (b) When both bidders pool, the second stage game has $\alpha = \beta = \frac{\theta\gamma}{1-\theta(1-\gamma)}$. (c) When one bidder pools but the other does not, the second stage game has $\beta = 1$ and $\alpha = \frac{\theta\gamma}{1-\theta(1-\gamma)}$. Using these second stage games, we compute the first stage profits from a non-pooling bid and a pooling bid:

$$\pi^{\text{CIP-semi},\theta}(p) = \theta(1 - F^{\text{CIP-semi},\theta}(p))(p - c_l)^{1-r} + (1-\theta)(c_h - c_l + p - c_l)^{1-r}. \tag{1}$$

$$\begin{aligned}\pi^{\text{CIP-semi},\theta}(c_h) &= \frac{1}{2}(c_h - c_l)^{1-r}((1-\theta)2^{1-r} + \theta\gamma) \\ &+\theta\int_{p_l}^{c_h}\left(\left(\frac{(1-\theta)(c_h - c_l + p - c_l)^{1-r}}{1-\theta(1-\gamma)} + \frac{\theta\gamma(p-c_l)^{1-r}}{1-\theta(1-\gamma)}\right)^{\frac{1}{1-r}} - (p - c_l)\right)^{1-r} dp \\ &+\frac{1-\theta+\theta\gamma}{2}\left(\left(\frac{(1-\theta)}{1-\theta+\theta\gamma}2^{1-r} + \frac{\theta\gamma}{1-\theta+\theta\gamma}\right)^{\frac{1}{1-r}} - 1\right)^{1-r}(c_h - c_l)^{1-r}\end{aligned} \tag{2}$$

We use these expressions to compute the first stage equilibrium in Appendix A.2.2.

Under CIP, our main focus is on the low-cost bidders' desire to pool in the first stage with the high-cost type. That incentive exists because, by preventing their opponent from learning about their type, bidders gain an information advantage for the second stage. To see this, consider bidder A who pools with the high-cost type in the first stage. Then, a Bayesian-updating bidder B will lower his belief in the second stage that bidder A is a low-cost type. This, in turn, allows bidder A to undercut bidder B with higher probability in the second stage. We refer to this pooling strategy in the first stage as the *deception effect*.

3.2.3 IIP Game

Under IIP, where only the winner's bid is revealed, the equilibrium in the first stage is separating in nature. Suppose $F^{\text{IIP},\theta}(p)$ and $f^{\text{IIP},\theta}(p)$ are the cdf and the pdf of the first stage bid distribution, and p_l^{IIP} represents the infimum of the bid distribution. Then, the second stage game in IIP has $\beta = 1$ and $\alpha = \frac{(1-F^{\text{IIP},\theta}(p_w))\theta}{1-\theta F^{\text{IIP},\theta}(p_w)}$, where p_w is the first stage winning bid. We account for the second stage game in computing the first stage profits:

$$\begin{aligned}\pi^{\mathrm{IIP},\theta}(p) &= \theta\int_{p_w=p_l^{IIP}}^{p_w=p}\left(\left(\alpha(c_h-c_l+p_w-c_l)^{1-r}+(1-\alpha)(p_w-c_l)^{1-r}\right)^{\frac{1}{1-r}}-(p_w-c_l)\right)^{1-r}f^{\mathrm{IIP},\theta}(p_w)dp_w \\ &\quad +(1-\theta F^{\mathrm{IIP},\theta}(p))(p-c_l)^{1-r}+(1-\theta)\left((c_h-c_l+p-c_l)^{1-r}-(p-c_l)^{1-r}\right) \qquad (3)\end{aligned}$$

We can use this equation to compute the first stage equilibrium (see Appendix A.3 for details.).

Note that there is no incentive under IIP for a low-cost bidder to pool with high-cost types in equilibrium. To see this, consider deviations by one of the bidders from the separating equilibrium identified under IIP. Instead of deviating and submitting a pooling bid, a low-cost bidder will be better off with a bid slightly less than the pooling bid. If the opponent is a high-cost type, a pooling bidder runs the risk of losing the first stage with a probability of one-half; with a bid slightly less than the pooling bid, she would have won the first stage with certainty. The pooling bidder also does not generate any benefit when facing a low-cost opponent. When competing against a low-cost opponent who is playing the strategy of the separating equilibrium, the pooling bid is never revealed. Hence, deviating to a pooling bid does not alter the first stage winner's belief for the second stage. The pooling bid does not allow the deviating bidder to learn about the opponent anymore than the bid slightly less than the pooling bid.

In IIP, only the *extraction effect* exhibited by the bidders comes into play. Since only the winner's bid is revealed under IIP, a low-cost bidder gains more information about his opponent for the second stage if he loses the first. A winner, however, gains less information. Thus, the policy creates an incentive for a low-cost bidder to bid higher and risk losing the first stage game to extract information about the opponent's type.

3.2.4 Analyses of Equilibrium Bids

This subsection characterizes properties of equilibrium bids that will be tested in the experiment. Kannan (2010) identifies the properties analytically under risk neutrality. In general, the equilibrium bid distributions are not analytically tractable for an arbitrary r. Hence, we characterize the equilibrium bid properties through numerical analysis for the θ values we use in our experiments, $\theta = \{0.5, 0.9\}$. These values satisfy the conditions in our theory for the extraction and the deception effects to exist.

Figures 1(a) to 1(d) show the first stage bid distributions under different CIP and IIP treatments

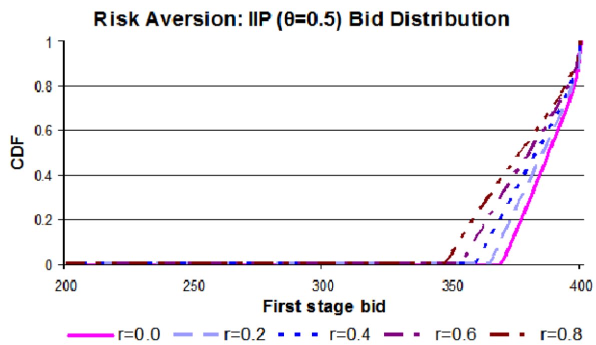

(a) IIP ($\theta = 0.5$)

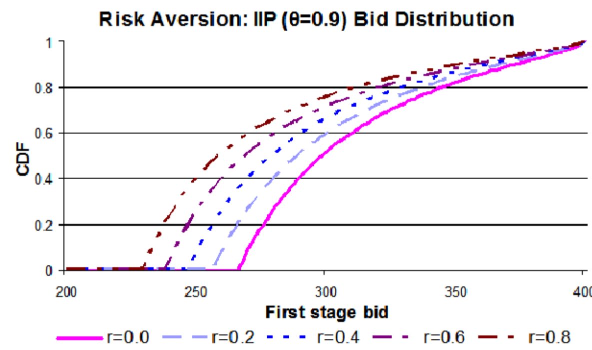

(b) IIP ($\theta = 0.9$)

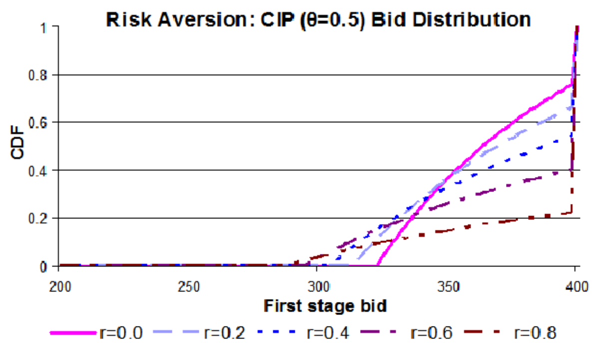

(c) CIP ($\theta = 0.5$)

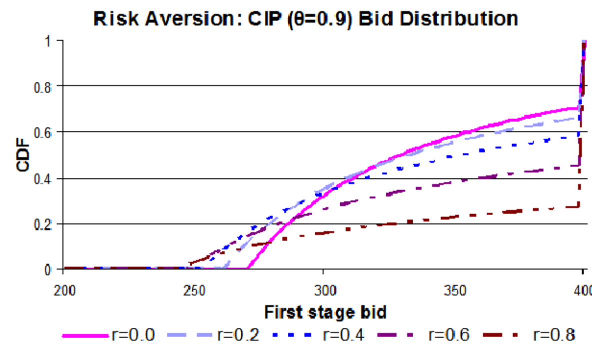

(d) CIP ($\theta = 0.9$)

Figure 1: First stage bid distributions for a low-cost type under IIP and CIP for different degrees of risk aversion when $c_l = 200$ and $c_h = 400$.

for various r values. Figures 1(a) and 1(b) illustrate that bid distributions shift leftwards under IIP as risk aversion increases. Hence, as risk aversion increases, individuals bid more aggressively to increase the probability of winning. The shape of the bid distributions also show that there is no positive probability for bids equal to c_h under IIP, i.e., no pooling bids are submitted under IIP. Comparing Figure 1(a) with Figure 1(b) shows that the bid distributions shift towards the left as θ increases. Hence, as the degree of competition increases, more weight is placed on lower bids.

Figures 1(c) and 1(d) show the bid distributions under CIP. Unlike for IIP, the bid distributions in CIP for the risk neutral case do not first-order stochastically dominate the risk averse cases. As risk aversion increases, the bid distributions shift leftwards only in the lower and intermediate bid ranges. In the upper bid range, however, the bid distributions shift downward. The downward shift in the bid distribution for bids of c_h illustrates that players submit more pooling bids as the degree of risk aversion increases. This result may be surprising at first glance since pooling increases the likelihood of losing in the first stage and therefore appears to be a risky strategy. However, in this multi-stage environment, one can show that pooling increases the likelihood of obtaining a positive payoff in the second stage and, hence, reduces the overall risk across the two stages. Recall that if both bids are smaller than c_h in the first stage, both bidders are revealed to be low cost and thus Bertrand competition results in zero profits in the second stage. Figure 2(a) shows the pooling rates under both CIP treatments for different r values. The pooling rates increase in the degree of risk aversion indicating the fact that bidders increase their pooling activity in both treatments. The pooling rates increase faster in the less competitive environment than in the more competitive environment.

Figure 2(b) shows the variation with respect to r of the expected first stage bid from low-cost bidders under CIP and IIP for the two θ values. Consider the learning effects when risk aversion (r) is low. The value gained from bid manipulations using both the extraction or the deception effect depends on the priors. The extraction effect corresponds to bidders learning if the opponent is a low-type. So, if the opponent is more likely to be a low-cost type anyway (i.e., θ is large), the additional information value a bidder acquires through extraction is low. The deception effect corresponds to a bidder preventing his opponent from learning about his low-cost type. So, if the opponent is expecting the bidder to be a high-cost type anyway (i.e., θ is small), the additional value from deceiving this opponent diminishes. Because of these two reasons, the extraction (deception) effect dominates for lower (higher) θ values, resulting in higher expected first stage bids under IIP (CIP), as observed in the figure.

As risk aversion increases, the comparison between CIP and IIP is also affected. The increasing pooling probability under CIP leads to higher expected prices under CIP. The bids in IIP ($\theta = 0.5$), however, decrease with an increase in r. Consequently, the difference between the first stage expected prices under CIP and IIP policies decrease with r. When $\theta = 0.5$, the dominance of the extraction effect over the deception effect can even vanish for very large levels of risk aversion (i.e., $r > 0.7$).

Figure 2(c) shows the expected price paid by the procurer across two stages in the four treatments. We find that when the competition is weak (strong), the prices paid in CIP are lower

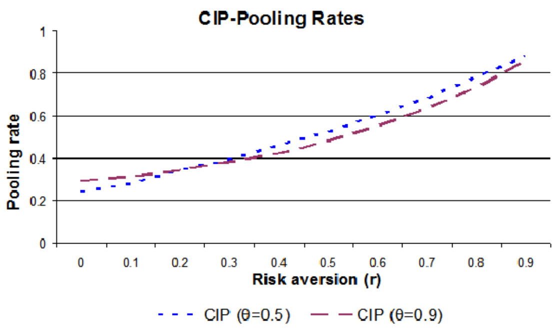

(a) Pooling rates

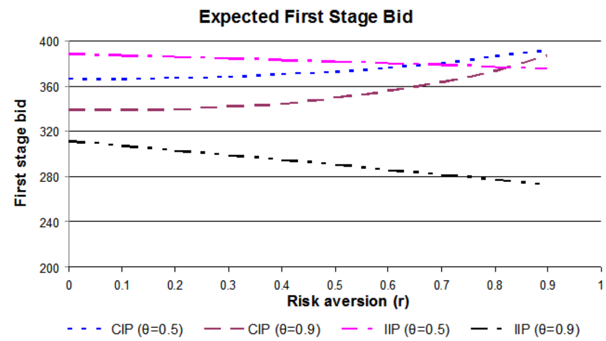

(b) Expected first stage prices

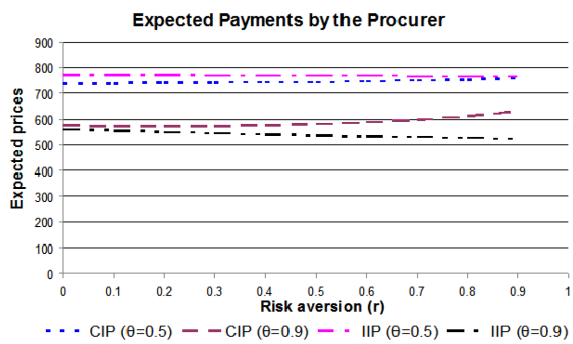

(c) Expected prices paid by the procurer

Figure 2: Pooling rates, expected first stage prices and expected price paid by the procurer for different treatments and different degrees of risk aversion.

(higher). The dominance of the extraction (deception) effect, which we noted earlier to occur when the competition is weak (strong), is also the reason for this observation. This is valid for any constant relative risk aversion parameter r.

4 Testable Hypotheses and Experiment Design

With the objective of experimentally studying the trade-offs between the two learning effects, the first subsection below builds on the equilibrium results in the previous section to develop hypotheses which we test in the experiment. The second section describes the experimental design.

4.1 Hypotheses

Note that the learning effects of interest are behaviors exhibited in the first stage by the low-cost types. Therefore, the first three of our four hypotheses are related to the first stage bids. The first hypothesis concerns pooling. For reasons discussed earlier, bidders submit pooling bids under CIP but not under IIP. Therefore,

Hypothesis 1 *The likelihood of pooling by a low-cost bidder is higher under CIP than under IIP.*

The second hypothesis is straightforward and intuitive from the equilibrium bid distributions (Figure 1).

Hypothesis 2 *As the probability of observing a low-cost opponent increases, the average price bid by the low-cost suppliers decreases.*

The third hypothesis involves the comparison of the extraction and the deception effects. Notice from Figure 2(b) that the dominance of one effect over the other depends on the risk aversion parameter. Prior work (e.g., Goeree et al., 2002, Campo et al., 2003, Holt and Laury, 2002) has identified a constant relative risk aversion parameter to be typically between $r = 0.3$ and $r = 0.6$. Corresponding to that range, our theory predicts higher average bids in CIP for the highly competitive ($\theta = 0.9$) treatment, and lower average bids for the less competitive ($\theta = 0.5$) treatment. Therefore,

Hypothesis 3 *When the probability of facing a low-cost opponent is low, the average price bid in the first stage by the low-cost suppliers is higher under IIP than CIP and vice-versa.*

While the first three hypotheses focus on the variations in the first stage bidding behavior, the next hypothesis concerns the procurer surplus over both stages. Based on the comparisons of the expected price paid by the procurer across both stages, we have the following hypothesis.

Hypothesis 4 *When the probability of facing a low-cost opponent is low, the total buyer payment across both periods is higher under IIP than CIP and vice-versa.*

Thus, we hypothesize that the procurer pays lower prices and obtains greater surplus on average by choosing CIP in less competitive conditions.

4.2 Experimental Design

To test these hypotheses, we employ an experimental framework that implements the stylized model summarized in Section 3 for the two values of θ (0.5 and 0.9). For the experiments, we set $c_l = 200$ and $c_h = 400$. From the theoretical results in Section 3.2.4, it should be clear that policy variations generate qualitative changes in bidder behavior because of the differences in how bidders process the information revealed. However, changes to θ merely generate quantitative variations in the strength of these incentives since the type of equilibrium remains unchanged. For this reason, we vary θ within the experimental sessions and vary the policies across sessions.

We created two datasets in the experiment, which differ in the restriction imposed on maximum bids (the details are provided later). Each experimental dataset comprises of 8 sessions, with four of them devoted to each information policy. Each session employs 12 subjects and involves 50 periods, and each period consists of two stages. An iid Bernoulli process is conducted in each period to determine the cost type for every subject, and this type remains the same for both stages of that period. Since we are testing a noncooperative equilibrium of a one-shot, two-stage game, we reduce repeated game incentives by randomly re-pairing the subjects with new opponents in each period. Although subjects interact repeatedly across the 50 periods, their interactions are anonymous and the identity of their interacting subject is never revealed. In each session, 25 periods are conducted with one θ treatment, followed by 25 periods for the other θ treatment. We vary the sequence of the θ treatments across sessions in order to control for possible order effects. Thus, our combined dataset consists of $9,600$ bids in each stage from 192 different participants across the 16 sessions.

For the first dataset, we assume that buyers have outside options to buy from suppliers if the bids exceed the high costs. For example, this could arise if the buyers can produce in-house at the same cost as the most inefficient potential supplier, allowing them to credibly commit to a reserve price of c_h (Thomas, 2010).[4] For this dataset we therefore explicitly restrict the maximum bid that subjects can submit to c_h. As we will present later, restricting the maximum bid to c_h is quite useful in providing managerial insights regarding the learning effects. The other advantage of this restriction is that it facilitates an accurate identification of pooling bids, which are defined to be equal to c_h. However, the restriction may appear to be too strong as it rules out certain strategies,

[4]This is similar to the standard restriction in the extensive experimental literature on buying auctions that bids are not allowed to be below the lowest possible buyer value.

such as collusion among potential high-cost suppliers. Furthermore, the imposed bidding cap may act as a focal point for bidders, leading to a second incentive for pooling besides the deception effect identified in the model for the CIP condition. Therefore, as a robustness check we also conduct a second set of experimental sessions.

This second dataset does not impose any explicit maximum bid restriction. We still, however, implicitly restrict the maximum bid to be $\approx \frac{35}{18}c_h$ (=777) since without a cap our exposure for potential monetary payments to subjects would be unbounded. Thus, in this design, the subjects know that a cap on the maximum bid exists but are not aware of its value. The restriction is not made explicit so that the revealed maximum bid does not become a focal point for the subjects. We also choose a non-intuitive limit that is large enough to facilitate potential collusion but also one that the subjects cannot easily arrive at through trial and error.[5] Thus, our second dataset overcomes the key potential problems with the first dataset. The one obvious disadvantage is that it becomes more challenging to define pooling bids, since high-cost types often bid above c_h. As we will discuss in the results section, this implies that low-cost bidders who submit bids of c_h might not be interpreted as pooling.

The experimental sessions were conducted in the Vernon Smith Experimental Laboratory at Purdue University. The subjects were recruited by email from the undergraduate student population and each subject was limited to participate in one session. Upon arrival at their experimental session, subjects were randomly assigned to individual computers and no communication between subjects was permitted throughout the session. At the start of the experiment, the instructions were read orally by an experimenter while the subjects followed along on their own copy. A sample of the instructions is provided in Appendix C. A computerized program written in z-Tree was used to implement the experimental environment (Fischbacher, 2007). The subjects in our experiment act as bidders, who receive monetary payments of their bid price minus their cost when they win an auction round and zero otherwise. The buyer's decision problem of picking the winner is computerized because buyers have no strategic role for exogenously-determined auction policies. All transactions and earnings are in experimental Francs, which are converted and paid in U.S. dollars at the end of experiment using a known and constant conversion rate. The feasible bids are in 0.01

[5]Subjects who bid above this threshold of 777 received the following "error" message from the experiment software: "The price you entered was too high. Please choose a lower price."

precision between 0 up to the maximum value imposed for that dataset.

In the theoretical analysis, bidders update their beliefs using Bayes' Rule. To understand the updating of beliefs in our experiment, we adopt a procedure (now common in the experimental economics literature) to elicit beliefs directly. In each period at the end of the first stage, every subject is asked to state – based on the information revealed – his belief (expressed as a probability) that his opponent has a cost of c_l. Monetary incentives are provided for making accurate guesses. Specifically, for eliciting the beliefs, we applied a quadratic scoring rule, which is incentive compatible for players to state their true their beliefs (Selten, 1998, Nyarko and Schotter, 2002, Costa-Gomes and Weizsacker, 2008). To reduce the likelihood that the belief elicitation reward significantly affects bidding behavior, the maximum reward for the beliefs (20 experimental Francs) is kept low relative to the difference in c_h and c_l.

The sequence of the experiment is as follows: Each bidder submits his bid for the first stage. The buyer buys from the lowest bidder (ties are broken randomly). After all bids are submitted, information about bids is revealed according to the chosen policy. Each subject then enters his belief about the opponent's cost type and his bid for the second stage. At the end of the second stage, the computers display the opponent's bids in the two stages and also the opponent's costs and the subject's own earnings. Subjects record all this information on hard copy record sheets so their personal history was easily accessible. Across all the 8-sessions, the earnings for the 192 subjects ranged from \$17.00 to \$40.00 with an average of \$25.25 per subject. Sessions typically lasted 90 to 100 minutes in total, including instruction time.

5 Results

In this section, we present the analyses of data obtained from our experimental sessions. Hereafter, we refer to the data collected from the first and the second set of experiments as Datasets 1 and 2, respectively. Table 1 provides the summary statistics of the bids submitted by the low-cost bidders for the first and the second stages for both datasets. The top panel is based on data from all 25 periods of each treatment sequence while the middle panel is only from the last fifteen periods in each treatment sequence. The first ten periods are excluded in the middle panel since those initial periods may involve subjects learning about the process of bidding. The bottom panel shows the theoretical mean values for the risk neutral case ($r = 0$). The variables used in the table correspond to the notation in the theoretical section: for a given treatment $l \in L = \{CIP, IIP\} \times \{\theta = 0.5, \theta = 0.9\}$, $\bar{p}^l$ is the average first stage bid submitted by the low-cost bidders, and $\bar{\gamma}^l$ is the rate with which a low-cost type bidder submits a bid ≥ 400.

Treatment		Dataset 1				Dataset 2			
		CIP		IIP		CIP		IIP	
$l \rightarrow$		$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$
All 25 periods included	$\bar{p}^l$	355.37	298.14	368.33	287.06	394.77	325.58	379.92	288.97
	Stdev	(44.94)	(60.02)	(41.52)	(44.42)	(43.40)	(74.48)	(50.61)	(52.65)
	$\bar{\gamma}^l$	8.16%	11.78%	2.21%	0.99%	49.66%	25.99%	32.31%	4.77%
	N	588	1,112	588	1,112	588	1,112	588	1,112
Only periods 11-25 included	$\bar{p}^l$	358.73	299.46	372.00	277.12	390.29	328.20	379.99	284.55
	Stdev	(39.87)	(64.76)	(35.38)	(41.47)	(31.34)	(77.85)	(40.76)	(49.24)
	$\bar{\gamma}^l$	4.83%	15.36%	1.14%	0.45%	44.03%	29.22%	28.69%	1.66%
	N	352	664	352	664	352	664	352	664
Theory $r = 0$	$\bar{p}^l$	366.08	339.33	385.77	309.97	366.08	339.33	385.77	309.97
	γ^l	23.95%	28.85%	0%	0%	23.95%	28.85%	0%	0%

Table 1: Descriptive statistics for bids submitted by low-cost bidders under CIP and IIP in both datasets.

One can note from the top two panels of the table that several observations are consistent with the theory across both datasets. For a given θ, the pooling probabilities are higher under CIP than under IIP in each of the datasets. The average bids in less competitive treatments ($\theta = 0.5$) are higher than the bids in more competitive ($\theta = 0.9$) treatments. Also, the bids are lower in IIP than in CIP when $\theta = 0.9$.

We also note that some results are not similar across the datasets (but those dissimilarities are valid across the top two panels). When $\theta = 0.5$, the average first stage bids from low-cost types are higher under CIP in Dataset 2 while the opposite holds in Dataset 1. Another result that is not consistent across the datasets is the pooling probability under the IIP ($\theta = 0.5$) treatment. For that treatment, while $\bar{\gamma}$ is close to zero and seems consistent with theory in Dataset 1, it is not the case in Dataset 2.

Table 1 shows that the average first stage bid in each dataset is lower for all but one treatment compared to the predicted Nash Equilibria under risk neutrality. The only exception occurs in the CIP ($\theta = 0.5$) treatment in Dataset 2. These observations can also be inferred from the bid distributions for all treatments shown in Figure 3, which usually indicate a greater weight on lower

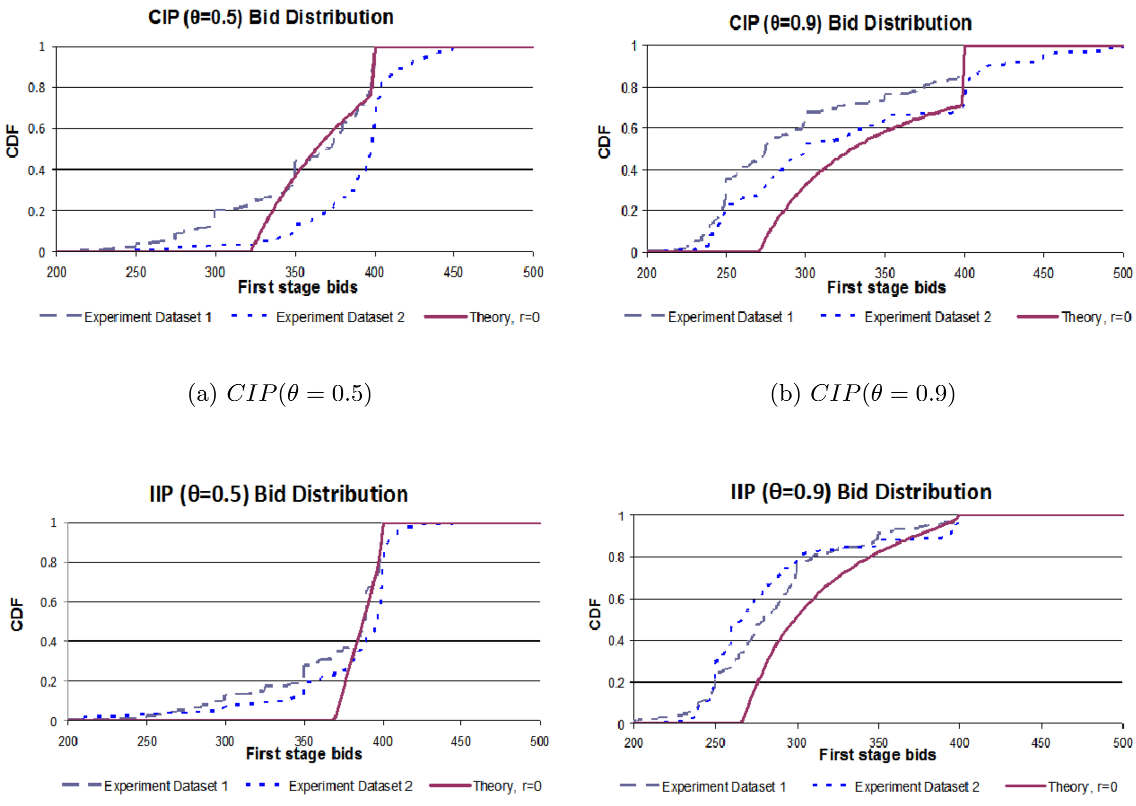

(a) $CIP(\theta = 0.5)$

(b) $CIP(\theta = 0.9)$

(c) $IIP(\theta = 0.5)$

(d) $IIP(\theta = 0.9)$

Figure 3: Theoretical and empirical first stage bid distributions for a low-cost type

bids. These increased frequencies on lower bids support the notion that risk aversion may help explain players' bidding behavior. We investigate the role of risk aversion later in Section 6.

5.1 Hypotheses Testing

The first three hypotheses concern the first stage bids from the low-cost sellers. We test these hypotheses using both datasets and apply regression models which include the following control variables: (i) the inverse of period (Inv_Period) to account for nonlinear time trends, and (ii) treatment order effects ($Treat\_Seq$), a dummy variable to differentiate the first and second treatments run within a session. These control variables allow for learning or other time series adjustments in behavior that are unrelated to the hypotheses of interest. In addition, we employ independent variables that are relevant for the respective hypothesis. The regressions for Hypothesis 1, 2, and 3 account for unobserved subject heterogeneity as random effects in order to control for additional

	Hypothesis 1				Hypothesis 2			
	Dataset 1 – Probit		Dataset 2 – Probit		Dataset 1 – Tobit		Dataset 2 – Tobit	
	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$	IIP	CIP	IIP	CIP
Dependent Var.	If a low-cost type pooled in the first stage				First stage bid from a low-cost type			
Intercept	-10.03** (1.18)	-4.78** (1.59)	-1.42** (0.57)	-3.46** (0.48)	362.91** (4.15)	358.29** (6.76)	384.31** (4.92)	402.81** (6.36)
Inv_Period	0.98 (0.89)	-1.97** (0.49)	1.59** (0.37)	-0.03 (0.25)	7.02* (3.56)	-21.02** (6.85)	5.61 (4.73)	-7.60 (7.07)
Treat_Seq	-0.37 (0.92)	0.28 (0.53)	-2.19** (0.64)	-0.33** (0.39)	18.94** (1.84)	17.30** (3.53)	-9.48** (2.43)	-17.35** (3.62)
Dummy_CIP	3.76** (1.11)	2.96** (1.34)	1.61** (0.63)	2.61** (0.48)				
Dummy_θ_0.9					-96.90** (1.84)	-58.14** (3.54)	-95.96** (2.44)	-63.98** (3.65)
Observations	704	1,328	704	1,328	1,016	1,016	1,016	1,016
Log L	-41.18	-193.67	-248.96	-329.12	-4,856.04	-4,988.29	-5,163.81	-5,562.71

Table 2: Regressions to test Hypotheses 1 and 2 using both datasets. Numbers in parentheses indicate the standard errors. Note that ** indicates a significance level of 1%, and * a significance level of 5%. Regressions include subject random effects. The first ten periods are omitted from the analysis.

factors not captured by our independent variables.

Hypothesis 1 concerns the impact of the policy on the propensity to submit pooling bids. We estimate a probit model with the dependent variable being a binary indicator that takes on a value of one when the low-cost bidder bids 400 in the first stage for both datasets. Since the focus here is on the information policy, we consider the bids from low-cost bidders corresponding to each θ separately but pool the data across the policies. The base case corresponds to IIP, and the dummy for CIP is the independent variable of interest. The regression coefficients are shown in the left panel of Table 2. They provide support for Hypothesis 1 for both datasets: the likelihood of observing pooling is significantly higher in CIP than IIP. Bidders perceive correctly that bidding to prevent an opponent from learning one's cost type is more useful under CIP than under IIP.[6]

For the regressions testing Hypothesis 2, the dependent variable is the bid submitted by a low-cost bidder. Since the maximum bid in Dataset 1 is 400, we employ a Tobit model with a 400 upper-bound. For Dataset 2, we employ a Tobit model as well with an upper limit of 777. We consider the bids submitted by low-cost bidders under each policy separately. The base case is when $\theta = 0.5$, and a dummy variable, *Dummy_θ_0.9*, is set to one for $\theta = 0.9$. The rightmost panel in Table 2 shows the regression results for both datasets. The coefficients on *Dummy_θ_0.9* are

[6] A Hausman test confirms that fixed effect results are not statistically different from the random effects results.

	Hypothesis 3				Hypothesis 4			
	Dataset 1 – Tobit		Dataset 2 – Tobit		Dataset 1 – Tobit		Dataset 2 – Tobit	
	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$
Dependent Var.	First stage bid from a low-cost type				Total price paid by the procurer			
Intercept	373.82**	262.87**	383.75**	292.67**	784.28**	497.75**	790.18**	525.76**
	(5.69)	(7.01)	(5.98)	(8.27)	(9.46)	(5.46)	(10.10)	(5.70)
Inv_Period	-5.92*	-7.51	1.58	-3.85	-28.65	-38.85**	-6.08	-24.48**
	(3.77)	(4.84)	(3.32)	(5.30)	(19.57)	(11.69)	(20.46)	(12.21)
Treat_Seq	4.28	31.73**	-10.78	-13.60	14.36	33.03**	-6.15	-34.60**
	(6.51)	(7.99)	(6.85)	(9.45)	(9.34)	(5.55)	(9.72)	(5.80)
Dummy_CIP	-14.87**	26.02**	10.46	43.30**	-38.99**	26.47**	-12.48	48.54**
	(6.51)	(7.99)	(6.85)	(9.45)	(9.37)	(5.55)	(9.75)	(5.80)
Observations	704	1,328	704	1,328	720	720	720	720
Log L	-3,252.38	-6,550.25	-3,273.51	-7,125.47	-2,668.32	-4,105.69	-2,473.01	-4,094.52

Table 3: Regressions to test Hypotheses 3 and 4 using both datasets. Numbers in parentheses indicate the standard errors. Note that ** indicates a significance level of 1% and * indicates significance level of 5%. Regressions for Hypothesis 3 include subject random effects. The first ten periods are omitted from the analysis.

negative and highly significant, consistent with Hypothesis 2. The data thus provide evidence that as the fraction of low-cost bidders increases, the first stage bids by low-cost bidders decrease.

To test Hypothesis 3, the regressions again use the bid submitted by a low-cost bidder as the dependent variable. We consider low-cost bids from each θ separately but pool the data across policies. The base case corresponds to IIP, and a dummy variable equal to one for CIP is the independent variable of interest. As before, we apply Tobit estimators for both datasets. The regression results are shown in the leftmost panel of Table 3. The results for Dataset 1 are consistent with the predictions of Hypothesis 3. The CIP dummy variable is negative and significant for the less competitive treatment, and positive and significant for the more competitive treatment. This indicates that the desire to learn dominates if the competition is weak (i.e., a lower θ value), so switching to a complete information policy lowers average bids. The results for Dataset 2, however, are in line with Hypothesis 3 only for the more competitive environment ($\theta = 0.9$). For the less competitive environment ($\theta = 0.5$) the estimate is not significantly different from zero. We investigate this deviation after testing Hypothesis 4.

For testing Hypothesis 4, the dependent variable is the sum of the prices paid by the buyer across the two stages. We adopt IIP as the base case and measure the impact of switching to CIP through the independent variable *Dummy_CIP*. As before, we use Tobit models for both datasets.[7] The

[7]Unlike the previous cases, these regression models do not include random subject effects because this market level performance measure depends on both sellers in each market and sellers are randomly re-paired each period.

regression results are shown in the rightmost panel of Table 3. For Dataset 1, the $Dummy\_CIP$ coefficient is negative and significant when $\theta = 0.5$, but positive and significant when $\theta = 0.9$. This result is consistent with Hypothesis 4, as the procurer pays less when adopting CIP only when the environment is less competitive. In more competitive environments, the procurer pays more when employing CIP since the low-cost bidders more frequently bid high prices to hide their type from the other bidder. As in the previous hypothesis, for Dataset 2, the result is consistent with our expectation only for $\theta = 0.9$. For $\theta = 0.5$, the $Dummy\_CIP$ coefficient is not significantly different from zero. This implies that a procurer pays about the same in CIP and IIP in an environment characterized by low competition. Note that, while much of the results are consistent with our expectations, the theoretically demonstrated dominance of the extraction effect over the deception effect for average bids and procurement costs is valid only in Dataset 1.

5.2 Pooling in Dataset 2

To understand better why the last two hypotheses receive only partial support in Dataset 2, consider again the pooling rates shown in Table 1 (Page 17). Note that these rates are significantly higher than the theoretical risk neutral case only for $\theta = 0.5$ in Dataset 2 under both policies. If we specifically focus on IIP across both datasets, observe from the table that the pooling rates are much higher than the theoretically predicted value of zero for $\theta = 0.5$ in Dataset 2, and are always higher in Dataset 2 than Dataset 1. These observations lead us to the following question: is the definition of a pooling bid, based on theoretical reasoning, appropriate for Dataset 2?

Recall that in equilibrium the high-cost type bids c_h regardless of their risk attitude, based on the standard reasoning behind Bertrand price competition. The maximum bid of c_h imposed in Dataset 1 is therefore not binding in equilibrium, but when this bid cap is relaxed in Dataset 2 many high-cost bidders bid greater than c_h. The bid distributions for these high-cost types shown in Figure 4 indicate that these higher bids are most common in the less competitive ($\theta = 0.5$) environment. Dufwenberg and Gneezy (2000) have shown in related experimental settings that prices are significantly higher than marginal costs in a two player Bertrand game. Similarly, our high-cost types sometimes bid prices higher than their marginal costs. This provides opportunities for the low-cost types to improve their profits with higher bids.

The price offers in Dataset 2 complicate our identification of a pooling bid by a low-cost type.

Nearly 88% of all bids submitted by high-cost types exceed $c_h = 400$, so by merely bidding at 400 a low-cost type is hardly pooling with the high-cost bidders behavior.

Pooling is only successful if it leads the other bidder to increase her belief that she is facing a high-cost type, so the belief data that we have elicited from bidders can provide direct evidence to infer pooling. This is easiest to illustrate in the CIP condition where all bids are observed. In Dataset 1, where bids above 400 are not possible, for $\theta = 0.5$ a bid equal to 400 in stage 1 causes the rival bidder to update her belief that the bidder is low cost to 0.17; and for $\theta = 0.9$ a bid equal to 400 causes the rival bidder to update her belief that the bidder is low cost to 0.39. Thus, a bid equal to 400 causes the rival to report a belief on average that the bidder is more likely to be a high-cost type.[8]

By contrast, in Dataset 2 for $\theta = 0.5$, the high-cost types bid strictly greater than 400 in stage 1 almost 92% of the time. Consequently, a bid equal to 400 causes the rival bidder to update her belief that the bidder is low cost to 0.57. That is, a bid of 400 leads the rival to believe that the high-cost type is *less* likely than the prior. For $\theta = 0.9$ the high-cost types bid strictly greater than 400 in stage 1 83% of the time, and a bid equal to 400 causes the rival bidder to update her belief that the bidder is low cost only to 0.70. Thus, in both cases a bid of 400 does not lead to a belief update that a bidder is substantially more likely to be high cost and cannot be considered a bid that is pooling with that type.

To pool with the high-cost types in Dataset 2, the low-cost types must bid somewhere within the distribution of bids chosen by the high-cost types, and not on the boundary of this distribution. Figure 4 shows that the distribution of these high-cost bids differs across treatments, with higher bids considerably more common when $\theta = 0.5$. Thus, an empirically-based definition of a threshold for identifying a pooling bid should vary with the treatment.

We considered a variety of alternatives for such an empirically-based definition of pooling bids. We settled on a straightforward definition of a pooling bid: One that is greater than or equal to the median bid submitted by the high-cost bidders in each treatment. Table 4 displays these treatment-specific median-bid thresholds.

This table indicates that the median bids from the high-type, h^{Median}, are approximately equal

[8]These beliefs should not be updated in equilibrium to indicate the high-cost type with certainty, of course, because some low-cost bidders are pooling with bids of 400.

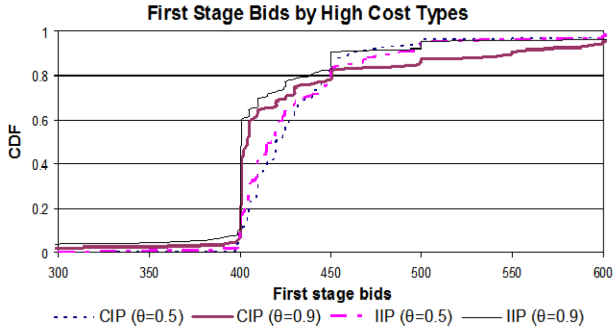

Figure 4: First stage bids by high-cost types, Dataset 2.

Treatment	CIP		IIP	
$l \rightarrow$	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$
h^{Median}	420	404	419	401
$\gamma_{p \geq h^{Median}}$	10.5%	16.1%	2.6%	0.6%

Table 4: Treatment-specific pooling bid thresholds in Dataset 2. h^{Median} is the median of bids from high-cost types, and $\gamma_{p \geq h^{Median}}$ is the probability that a low-cost type bids equal to or above h^{Median}.

between the CIP ($\theta = 0.5$) and IIP ($\theta = 0.5$) treatments. However, the frequency that the bids submitted by a low-cost type are greater than or equal to h^{Median} is much higher under CIP than IIP. This is consistent with the theoretical prediction that the deception effect is only present in the CIP condition. Regardless of whether the pooling definition is adjusted as just described, pooling is more common in CIP (Table 1), and regression results (available upon request) using the median price for the pooling definition provide conclusions that are similar to those drawn from Table 2.

This complication for Dataset 2 arises because of non-Bertrand behavior by high-cost types in stage 1. This may be because only two bidders compete in each auction. Dufwenberg and Gneezy (2000) showed that Bertrand behavior is much more common with three or four sellers, which leads us to conjecture that the high-cost types will typically submit bid equal to their marginal costs when more than two bidders compete. Thus, differences between Datasets 1 and 2 arising from different bid caps may decrease substantially in less concentrated markets, with both behaving similarly to Dataset 1 with its c_h cap. Hence, we believe that the insights from the analysis of Dataset 1 are more broadly relevant.

6 Risk Aversion

The previous section shows that the experimental data are broadly consistent with the comparative static predictions of the theoretical model. Subjects exhibit behavior that appears to appreciate the two learning effects. However, several quantitative deviations exist. As pointed out earlier, most notably the observed average bids in Table 1 (Page 17) are lower for most treatments than the risk neutral theoretical prediction. We next investigate how well risk aversion accounts for this deviation. Note that risk aversion has also been proposed as an explanation for such deviations in non-experimental auction data, even with large firms as bidders (e.g., Campo et al., 2003). In this section, we use the theoretical results from Section 3 and follow a structural approach to estimate the degree of risk aversion that is most consistent with observed behavior.

While the process can be applied to both our datasets, we only present the results for Dataset 1. In Dataset 2, the estimates obtained through the same procedure are problematic because, as documented above, the bids from the high-cost types do not conform to the Bertrand prediction of marginal cost offers. There is no easy structural way of accounting for this deviation, since it leads low-cost bidders to bid above c_h, which never occurs in equilibrium for any level of risk aversion. For completeness, however, we report the risk aversion estimations for Dataset 2 in Appendix B.

6.1 Structural Approach for Estimating Risk Aversion

In this section, we account for mixed equilibrium strategies and apply a structural approach that relies on the assumption that observed bids are generated from the equilibrium model, allowing for risk aversion. Specifically, we estimate the degree of risk aversion that minimizes the distance between the theoretical and the empirical bid distributions.

Let $p_1^l, \ldots, p_n^l$ be the observed first stage bids from low-cost suppliers under treatment $l \in L = \{CIP, IIP\} \times \{\theta = 0.5, \theta = 0.9\}$,[9] and X^l be a set such that all the observed first stage bids from low-cost types are included distinctly (only once). The cdf of the observed bid distribution for an arbitrary bid x is $\frac{1}{n}\sum_{i=1}^{n} I\left(p_i^l \leq x\right) \equiv \tilde{F}^l(x)$, where I is the indicator function. Let the bids originate from an underlying distribution with cdf for x being $F^l(x, \widehat{r})$, where $\widehat{r}$ is the risk aversion coefficient of the subjects. For any arbitrary risk coefficient r, $F^l(x, r)$ indicates the cdf of the theoretically generated bid distribution. For any x, the squared difference between the cdfs of the

[9] Note that $\bar{p}^l$ in Table 1 is $\frac{1}{n}\sum_{i=1}^{n} p_i^l$.

theoretical distribution computed at a risk aversion coefficient of r and the empirical distribution is given by

$$\epsilon_r^l(x) \quad = \quad \left(\tilde{F}^l(x) - F^l(x,r)\right)^2 .$$

The objective is to find the risk aversion parameter that minimizes the sum of squared differences between the theoretical and empirical bid distributions:[10]

$$\widehat{r} \quad = \quad \arg\min_r \sum_{l \in L} \frac{1}{|X^l|} \sum_{x \in X^l} \epsilon_r^l(x). \tag{4}$$

The differences $\epsilon_r^l(x)$ can be interpreted as the errors arising from the agent's optimization, or from considering a subset of the whole population. Since no closed form solution exists for the equilibrium with non-neutral risk preferences, we apply a grid search over r in increments of 0.01.

In order to test for statistical significance as well as equality of risk aversion estimates between treatments we compute the standard errors using the non-parametric bootstrap method (with replacement) as suggested by Efron (1982). Our estimate of the standard error is obtained using 500 bootstrap samples from the empirical distribution of the data.[11]

6.2 Results from the Estimation Procedure

We estimate the risk aversion coefficient using the datasets with the first 10 periods in each treatment omitted in order to reduce the influence of learning effects and adjustments to the treatment conditions.[12] The risk aversion coefficient that best fits the $2,032$ observations pooled across all treatments is $r = 0.38$ and the bootstrapped standard error is 0.03. Thus, the estimate is significantly different from zero (risk neutrality). This estimate is consistent with previous findings and well-documented in experimental research focusing on first price auctions with competing buyers; see for instance Goeree et al. (2002), Cox et al. (1988), Harrison (1989) and Kagel (1995).

[10] In order to check for the robustness of this metric, we also use tests based on the supremum distance norm as suggested by Romano (1988) and Romano (1989). We use a Kolmogorov-Smirnov test using the supremum distance between the theoretical and empirical bid distributions and search over the different risk aversion parameters according to the minimum distance principle. The optimal risk aversion parameter satisfies $\widehat{r} = \arg\min_r \sup \left|\tilde{F}^l(x) - F^l(x,r)\right|$. The results are quite similar for most of the treatments, and they are available from the authors upon request.

[11] For tests at level 0.05, Efron and Tibsharani (1993) and Davidson and MacKinnon (2000) recommend 200 and 399 samples, respectively.

[12] This is the same subsetting to later periods employed in the initial equilibrium hypothesis testing in Section 5. The estimation results using all periods are similar to the reported ones.

		All Treatments	CIP		IIP	
			$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$
Dataset 1	Risk aversion coefficient	0.38**	0.01	0.26**	0.57**	0.68**
	Stdev	(0.03)	(0.01)	(0.02)	(0.17)	(0.02)
	Observations	2,032	352	664	352	664

Table 5: Results for estimated risk aversion coefficients. Numbers in parentheses indicate the bootstrapped standard errors. Note that ** indicates a significance level of 1%. The first ten periods are omitted from the estimations.

Table 5 shows the risk aversion coefficient estimated overall and separately for every treatment. Interestingly, the only treatment in which the data appear most consistent with risk neutrality ($r = 0$) is the CIP ($\theta = 0.5$) treatment, which we investigate further.

For this CIP ($\theta = 0.5$) treatment, Table 6 shows the fit error for various bid price ranges based on the treatment-specific risk estimate (which corresponds to risk neutrality, $r = 0$) and the overall risk estimate (which is $r = 0.38$). The overall risk estimate appears to be a better fit for lower bid prices. It is only in the higher bid prices does the fit become better with the risk neutral case. One of the key factors contributing to the better fit at higher bid prices is the pooling rate. Notice that the observed pooling rate (0.05 in these late periods) is much lower than the risk neutral equilibrium (0.24), and is smaller than any higher theoretical risk aversion rate (cf Figure 2(a)). Consequently, risk aversion does not fit better in the last two rows in Table 6. The aggressive bidding in general seems to indicate bidder behavior that is consistent with risk aversion. However, the lower pooling rate and the continued aggressive bidding behavior at higher bid prices suggests that bidders do not recognize the somewhat counter-intuitive incentive to increase their pooling rate with risk aversion. This may be attributed to the cognitive limitation of our human subjects. While we have only focused on risk aversion, other motivations such as a non-monetary utility of winning the auction (Sheremeta, 2010), regret (Engelbrecht-Wiggans and Katok, 2008) or spite/envy Morgan et al. (2003), can also influence bidder behavior.

7 Conclusion

This study represents a first step to provide empirical evidence regarding repeated auctions with information asymmetry and independent private values. We experimentally investigate how different information revelation policies affect submitted bids in sequential auctions with private information.

χ	$\sum_\chi \epsilon_{r=0}^{\text{CIP},0.5}$	$\sum_\chi \epsilon_{r=0.38}^{\text{CIP},0.5}$
≤ 300	0.056	0.056
≤ 320	0.179	0.092
≤ 340	0.225	0.096
≤ 360	0.237	0.117
≤ 380	0.258	0.216
≤ 390	0.264	0.412
≤ 400	0.385	1.018

Table 6: Sum of squared errors under CIP, based on Dataset 1.

Motivated by practical concerns faced by managers involved in intermediate goods transactions, our analysis considers two stage procurement auctions in which bidders are uncertain about their opponent's cost structure. We compare two different information revelation policies that could be chosen by the procurer, one in which only the winning bid is revealed (IIP) and one in which all bids are revealed (CIP).

Our analysis also provides several important managerial insights for participants in these traditional and online procurement markets. It shows that human bidders, even with limited training, indeed appreciate the intuitions regarding the extraction and the deception effects. The theoretical and empirical analysis shows that either effect can dominate the other, so buyer surplus could be higher in either auction format depending on the perceived competitiveness in the market. Therefore, both these effects are important and should be considered when choosing the appropriate information revelation policy. Although the theoretical model with risk averse bidders shows a surprising result – the direct variation of the pooling probability with the degree of risk aversion – the experimental subjects do not appear to correctly appreciate the risk reduction provided by pooling. They bid aggressively even when pooling could reduce risk, perhaps due to an extra utility of winning. We recommend that a manager also consider such cognitive limitations and winning incentive when choosing the auction rules.

One may wonder about the applicability of our results beyond the two cost types studied. As Lambson and Thurston (2006) note, a model assuming discrete type bidders is more realistic than continuous type bidders. The discreteness naturally leads to bidders exhibiting learning-related behaviors (Jeitschko, 1998), which is the focus of our paper. For an arbitrary number of consumer types, the analyses of the two policies are not tractable. For the three-cost type framework, Kannan (2010) has already shown the existence of the extraction and the deception effects. Using numerical analyses, he demonstrates cases when one effect dominates the other. Given how remarkably subjects in our experiment appreciate the learning behaviors without even computing the equilibria, we believe that those learning behaviors will be relevant even to settings with an arbitrary number of cost types.

Future research could investigate whether the model is also generally well-supported if bidders face more than one opponent. As theoretically shown by Kannan (2010), when the number of bidders increases, this changes the equilibrium bidding behavior in the first and the second stages. It would also be useful to broaden the economic environment to include richer, possibly continuous cost distributions, as well as more than two periods. Although these extensions have not been addressed theoretically for repeated auctions of this type due to intractability, it is straightforward to extend the present laboratory setup in this direction to assess the robustness of our main empirical conclusions.

References

Anand, K. and M. Goyal (2009). Strategic Information Management under Leakage in a Supply Chain. *Management Science 55*(3), 438–452.

Arora, A., A. Greenwald, K. N. Kannan, and R. Krishnan (2007). Effects of Information-Revelation Policies Under Market-Structure Uncertainty. *Management Science 53*(8), 1234–1248.

Bajari, P. (2001). Comparing Competition and Collusion in Procurement Auctions:A Numerical Approach. *Economic Theory 18*, 187–205.

Branco, F. (1997). The Design of Multidimensional Auctions. *Rand Journal of Economics 28*, 63–81.

Camerer, C. and T.-H. Ho (1999). Experience-weighted Attraction Learning in Normal Form Games. *Econometrica 67*(4), 827–874.

Campo, S., E. Guerre, I. Perrigne, and Q. Vuong (2003). Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders. CREST working paper.

Chandrashekar, T., Y. Narahari, C. Rosa, D. Kulkarni, J. D. Tew, and P. Dayama (2007). Auction-Based Mechanisms for Electronic Procurement. *IEEE Transactions on Automation Science and Engineering 4*(3), 297–321.

Chen, Y. and G. Vulcano (2009). Effects of Information Disclosure Under First- and Second-Price Auctions in a Supply Chain Setting. *Manufacturing and Service Operations Management 11*(2), 299–316.

Costa-Gomes, M. and G. Weizsacker (2008). Stated Beliefs and Play in Normal-Form Games. *Review of Economic Studies 75*, 729–762.

Cox, J., V. Smith, and J. Walker (1988). Theory and Individual Behavior of First-Price Auctions. *Journal of Risk and Uncertainty 1*, 61–100.

Das Varma, G. (2003). Bidding for a Process Innovation Under Alternative Modes of Competition. *International Journal of Industrial Organization 21*, 15–37.

Davidson, R. and J. MacKinnon (2000). Bootstrap Tests: How Many Bootstraps? *Econometric Review 19*, 55–68.

de Silva, D., T. Dunne, A. Kankanamge, and G. Kosmopoulou (2008). The Impact of Public Information on Bidding in Highway Procurement Auctions. *European Economic Review 52*, 150–181.

Donald, S. G., H. J. Paarsch, and J. Robert (2006). An Empirical Model the Multi-Unit, Sequential Clock Auction. *Journal of Applied Econometrics 21*, 1221–1247.

Dufwenberg, M. and U. Gneezy (2000). Price Competition and Market Concentration: An Experimental Study. *International Journal of Industrial Organization 18*, 7–22.

Dufwenberg, M. and U. Gneezy (2002). Information Disclosure inAuctions: An Experiment. *Journal of Economic Behavior and Organization 48*(4), 431–444.

Efron, B. (1982). The Bootstrap and other Resampling Plans. Society for Industrial and Applied Mathematics.

Efron, B. and J. Tibsharani (1993). *An Introduction to the Bootstrap*. London: Chapman and Hall.

Elmaghraby, W. and P. Keskinocak (2000). Technology for Transportation Bidding at the Home Depot. Teaching case available from `http://www2.isye.gatech.edu/people/faculty/Pinar_Keskinocak/home-depot-case.zip`.

Engelbrecht-Wiggans, R. and E. Katok (2008). Regret and Feedback Information in First-Price Sealed-Bid Auctions. *Management Science 54*(4), 808–819.

Englebrecht-Wiggans, R. (1994). Sequential Auctions of Stochastically Equivalent Objects. *Economics Letters 44*, 87–90.

Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments,. *Experimental Economics 10*(2), 171–178.

Goeree, J. K. (2003). Bidding for the Future: signaling in Auctions with an Aftermarket. *Journal of Economic Theory 108*, 345–364.

Goeree, J. K., C. A. Holt, and J. O. Ledyard (2006). An Experimental Comparison of the FCCs Combinatorial and Non-Combinatorial Simultaneous Multiple Round Auctions. FCC, Working Paper.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2002). Quantal Response Equilibrium and Overbidding in Private-Value Auctions. *Journal of Economic Theory 104*(1), 247–272.

Harrison, G. (1989, September). Theory and Misbehavior of First-Price Auctions. *American Economic Review 79*, 749–762.

Hausch, D. (1986). Multi-object Auctions: Sequential vs Simultaneous sales. *Management Science 32*(12), 1599–1610.

Holt, C. and S. Laury (2002, January). Risk Aversion and Incentive Effects. *American Economic Review 92*, 1644–1655.

Jap, S. D. (2002). Online Reverse Auctions: Issues, Themes, and Prospects for the Future. *Journal of the Academy of Marketing Science 30*, 506–25.

Jap, S. D. (2003, July). An Exploratory Study of the Introduction of Online Reverse Auctions. *Journal of Marketing 67*, 96–107.

Jeitschko, T. (1998). Learning in Sequential Auctions. *Southern Economic Journal 65*(1), 98–112.

Jeitschko, T. and Wolfstetter (1998). Scale Economies and the Dynamics of Recurring Auctions. *Economic Inquiry 40*(3), 403–414.

Kagel, J. (1995). Auctions: A Survey of Experimental Research. In J. H. Kagel and A. E. Roth (Eds.), *The Handbook of Experimental Economics*, Princeton, New Jersey. Princeton University Press.

Kagel, J. and D. Levin (2010). Auctions: A Survey of Experimental Research, 1995-2008. In J. Kagel and A. Roth (Eds.), *Handbook of Experimental Economics*, Volume 2, Princeton, New Jersey. Princeton University Press.

Kannan, K. (2010). Effect of Information Revelation Policies Under Cost Structure Uncertainty. Forthcoming in *Information Systems Research.*

Katzman, B. and M. Rhodes-Krop (2008). The Consequences of Information Revealed in Auctions. *Applied Economics Research Bulletin 2*, 53–87.

Klemperer, P. (2004). *Auctions: Theory and Practice.* Princeton, New Jersey: Princeton University Press.

Krishna, V. (2009). *Auction Theory.* San Diego, California: Academic Press.

Lambson, V. and N. Thurston (2006). Sequential auctions: Theory and Evidence from the Seattle Fur Exchange. *RAND Journal of Economics 37*(1), 70–80.

Maskin, E. and J. Riley (1984). Optimal Auctions with Risk Averse Buyers. *Econometrica 52*(6), 1473–1518.

Maskin, E. and J. Riley (1985). Auction Theory in Private Values. *American Economic Review 75*(2), 150–155.

Maskin, E. and J. Riley (1987). Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica 55*, 633–646.

Maskin, E. and J. Riley (2000). Equilibrium in Sealed High Bid Auctions. *The Review of Economic Studies 67*(3), 439–454.

Milgrom, P. R. and R. J. Weber (1982). A Theory of Auctions and Competitive Bidding. *Econometrica 50*, 1089–1122.

Morgan, J., K. Steiglitz, and G. Reis (2003). The Spite Motive and Equilibrium Behavior in Auctions. *Contributions to Economic Analysis and Policy 2*(1), Article 5.

Nyarko, Y. and A. Schotter (2002). An Experimental Study of BeliefLearning Using Elicited Beliefs. *Econometrica 70*(3), 971–1005.

Ortega-Reichert, A. (1968). *Models for Competitive Bidding Under Uncertainty.* Ph. D. thesis, Department of Operations Research, Stanford University.

Romano, J. (1988). A Bootstrap Revival of Some Nonparametric Distance Tests. *Journal of the American Statistical Association 83*, 698–708.

Romano, J. (1989). Bootstrap and Randomization Tests of Some Nonparametric Hypotheses. *Annals of Statistics 17*, 141–159.

Selten, R. (1998). Axiomatic Characterization of the Quadratic Scoring Rule. *Experimental Economics 1*, 43–61(19).

Sheremeta, R. (2010). Experimental Comparison of Multi-Stage and One-Stage Contests. *Games and Economic Behavior 68*, 731–747.

Thomas, C. (1996). Market Structure and the Flow of Information in Repeated Auctions. Working paper, Federal Trade Commission.

Thomas, C. (2010). Information Revelation and Buyer Profits in Repeated Procurement Competition.

Journal of Industrial Economics 58(1), 79–105. Journal of Industrial Economics.

Tu, Z. (2005). Why do we use Dutch Auction to Sell Flowers? Information Disclosure in Sequential Auctions. Working Paper, Department of Economics, University of Pittsburgh.

Wakker, P. (2008). Explaining the Characteristics of the Power (CRRA) Utility Function. *Health Economics 17*, 1329–1344.

Appendix A: Derivation of the Theoretical Model allowing for Risk Aversion

A.1 Second Stage Game Equilibrium

Note that the bid price infinitesimally smaller than c_h is the supremum of the strategy space and, in a mixed strategy equilibrium, payoffs from any price q should yield the same payoff. Based on that, we can compute

$$F_B(q) = \frac{1}{\alpha}\left(1-(1-\alpha)\frac{(c_h-c_l+p-c_l)^{1-r}-(p-c_l)^{1-r}}{(p-c_l+q-c_l)^{1-r}-(p-c_l)^{1-r}}\right).$$

We find that the infimum of $F_B(q)$ is $\underline{q} = 2c_l - p + ((1-\alpha)(c_h-c_l+p-c_l)^{1-r}+\alpha(p-c_l)^{1-r})^{\frac{1}{1-r}}$. This price will also correspond to the infimum of $F_A(q)$. Using that information, we compute

$$F_A(q) = \frac{1}{\beta}\left(1-(\frac{\underline{q}-c_l}{q-c_l})^{1-r}\right)$$

and a masspoint of $\frac{1}{\beta}(\frac{\underline{q}-c_l}{q-c_l})^{1-r}$ at the price infinitesimally smaller than c_h. In this case, the second stage expected profits for A is $(1-\alpha)((c_h-c_l+p-c_l)^{1-r}-(p-c_l)^{1-r})$, and for B is $(\underline{q}-c_l)^{1-r}$. Using these cdfs, one can directly obtain the pdf of the bid distributions. Let the respective probability density functions (pdfs) be $f_A(q)$ and $f_B(q)$. Note that even if the beliefs are symmetric, the bid distributions may be different because of A's gain from the previous period. Also, by rearranging $\pi_A(q)$, we obtain the following expression, which will be useful later:

$$(1-\alpha F_B(q))\left((p-c_l+q-c_l)^{1-r}-(p-c_l)^{1-r}\right) = (1-\alpha)((c_h-c_l+p-c_l)^{1-r}-(p-c_l)^{1-r}) \quad (5)$$

A.2 First Stage Game: CIP

Recall that under CIP, all bids are revealed. Two cases are possible

A.2.1 Separating Equilibrium Case

Note that the price infinitesimally smaller than c_h is in the supremum of the strategy set and the expected profit from any p in the mixed strategy equilibrium strategy set should be the same. Based on that, we can obtain the expected profit from any p is $(2(c_h - c_l))^{1-r}(1-\theta)$ and $F^{\text{CIP-sep},\theta}(p) = 1 - \frac{1-\theta}{\theta}\left(\frac{(2(c_h-c_l))^{1-r}-(c_h-c_l+p-c_l)^{1-r}}{(p-c_l)^{1-r}}\right)$. As expected, the bids tend to be lower with increasing risk aversion (=r) and higher degrees of competition (θ). The separating equilibrium is only sustained under a certain condition. To determine when the separating equilibrium exists, consider deviations by one of the bidders. Suppose a bidder pretends to be a high-cost type in the first stage. The expected profit for that bidder is:

$$\pi^{pool} = \overbrace{\overbrace{(1-\theta)\frac{1}{2}(2(c_h-c_l))^{1-r}}^{\text{Bidder wins both stages}} + \overbrace{(1-\theta)\frac{1}{2}(c_h-c_l)^{1-r}}^{\text{Bidder wins second stage}}}^{\text{High-cost rival}} + \overbrace{\overbrace{\theta(c_h-c_l)^{1-r}}^{\text{Bidder wins second stage}}}^{\text{Low-cost rival}}.$$

We can observe that $\pi^{pool} > (2(c_h-c_l))^{1-r}(1-\theta)$, i.e., pooling is more profitable, if $\theta > \frac{2^{1-r}-1}{2^{1-r}+1}$.

A.2.2 Semipooling Equilibrium Case

We first define the variables for the second stage game. For case (b), the cdf of the first winner's bid distribution is denoted by $F_{ws}(q)$ and of the loser by $F_{ls}(q)$; the subscript s corresponds to the case with symmetric beliefs. For case (c), the winner bids according to the cdf $F_{wa}(q)$, and the loser according to $F_{la}(q)$, where the subscript a represents the asymmetric case. The corresponding pdfs are $f_{ws}(q)$, $f_{ls}(q)$, $f_{wa}(q)$ and $f_{la}(q)$.

We next characterize the expected payoffs under two scenarios: (i) when the bid in the first stage reveals the type, and (ii) when the low-cost bidder pools with a high-cost type. Consider scenario (i). The profit from bidding $p < c_h$ in the first stage and q in the second is:

$$\begin{aligned}
\pi^{\text{CIP-semi},\theta}(p,q) = \; & \overbrace{(1-\theta)}^{\text{High-cost rival}} \; f^{\text{CIP-semi},\theta}(p) f_{wa}(q)(p-c_l+q-c_l)^{1-r} \\
& + \overbrace{\theta(1-F^{\text{CIP-semi},\theta}(p)-\gamma)}^{\text{Bidder wins stage 1; Both low-costs revealed} \implies \text{0 profit in stage 2}} \; f^{\text{CIP-semi},\theta}(p)(p-c_l)^{1-r} \\
& + \overbrace{\theta\gamma(1-F_{la}(q))}^{\text{Bidder wins both stages; low-cost rival hides}} \; f^{\text{CIP-semi},\theta}(p) f_{wa}(q)(p-c_l+q-c_l)^{1-r} \\
& + \overbrace{\theta\gamma F_{la}(q)}^{\text{Hiding low-cost rival wins stage 2}} \; f^{\text{CIP-semi},\theta}(p) f_{wa}(q)(p-c_l)^{1-r}
\end{aligned}$$

By rearranging and integrating with respect to the second stage bid, we obtain the conditional profit of bidding p as:

$$\begin{aligned}\pi^{\text{CIP-semi},\theta}(p) &= (1-\theta F^{\text{CIP-semi},\theta}(p))(p-c_l)^{1-r} + \\ &\quad \int_{q=\underline{q}}^{q=c_h}(1-\theta+\theta\gamma(1-F_{la}(q)))\left((p-c_l+q-c_l)^{1-r}-(p-c_l)^{1-r}\right)f_{wa}(q)dq \\ &= (1-\theta F^{\text{CIP-semi},\theta}(p))(p-c_l)^{1-r}+(1-\theta)((c_h-c_l+p-c_l)^{1-r}-(p-c_l)^{1-r})\end{aligned}$$

The second equality is due to Equation 5. Simplifying the expression, we get 1. The expected profit from any bid p can be obtained by setting $F^{\text{CIP-semi},\theta}(p)$ for a p price infinitesimally smaller than c_h to be $1-\gamma$. Therefore,

$$\pi^{\text{CIP-semi},\theta}(p) = (\theta\gamma+(1-\theta)2^{1-r})(c_h-c_l)^{1-r}. \tag{6}$$

Next, we focus on scenario (ii). Consider the profit from submitting a pooling bid when q is bid in the second stage:

$$\begin{aligned}\pi^{\text{CIP-semi},\theta}(c_h,q) &= \overbrace{\frac{\overbrace{(1-\theta)}^{\text{High-cost rival}}+\overbrace{\theta\gamma(1-F_{ls}(q))}^{\text{Hiding low-cost rival}}}{2}(c_h-c_l+q-c_l)^{1-r}f_{ws}(q)}^{\text{Bidder wins both stages}} \\ &\quad + \overbrace{\frac{\theta\gamma F_{ls}(q)}{2}(c_h-c_l)^{1-r}f_{ws}(q)}^{\text{Hiding low-cost rival and bidder wins stage 1}} + \overbrace{\frac{\overbrace{(1-\theta)}^{\text{High-cost rival}}+\overbrace{\theta\gamma(1-F_{ws}(q))}^{\text{Hiding low-cost rival}}}{2}(q-c_l)^{1-r}f_{ls}(q)}^{\text{Bidder wins only second stage}} \\ &\quad \overbrace{+\theta\int_{p_w=p_l^{\text{CIP}}}^{p_w=c_h}\left((1-F_{wa}(q))(q-c_l)^{1-r}f_{la}(q)\right)f^{\text{CIP-semi},\theta}(p_w)dp_w}^{\text{Revealing low-cost rival and bidder wins stage 2}}\end{aligned}$$

Corresponding to the first three terms, the winning bid price in the first stage is always c_h, and for the last one it is some arbitrary $p < c_h$. The first two terms together relate to Π_A, and the third and the fourth terms relate to Π_B in the second stage game. Using them, we have Equation 2.

We solve for the two variables $F^{\text{CIP-semi},\theta}(p)$ and γ using two equations, which are obtained by setting any two of the following three equations equal: expected profit in Equation 1 for any $p < c_h$; the expected profit computed at the supremum of the strategy set, which is infinitesimally smaller than c_h, in Equation 6; and Equation 2. The equilibrium bid distribution does not have a closed form expression for any arbitrary r but is available for $r = 0$ (Kannan, 2010).

A.3 First Stage Game: IIP

We again begin with the definition of the second stage variables. We represent $F_w(q)$ and $f_w(q)$ as those for the first stage winner in the second stage; and $F_l(q)$ and $f_l(q)$ for the first stage loser in the second stage. The expected payoff from bidding p in the first stage and q in the second stage is

$$\begin{aligned}
\pi^{\mathrm{IIP},\theta}(p,q) &= \overbrace{(1-\theta)f_w(q)(p-c_l+q-c_l)^{1-r}}^{\text{High-cost rival}} f^{\mathrm{IIP},\theta}(p) \\
&\quad + \overbrace{\theta(1-F^{\mathrm{IIP},\theta}(p))(1-F_l(q))f_w(q)(p-c_l+q-c_l)^{1-r}}^{\text{Low-cost rival exists but bidder wins both stages}} f^{\mathrm{IIP},\theta}(p) \\
&\quad + \overbrace{\theta(1-F^{\mathrm{IIP},\theta}(p))F_l(q)f_w(q)(p-c_l)^{1-r}}^{\text{Low-cost rival exists but bidder wins stage 1}} f^{\mathrm{IIP},\theta}(p) \\
&\quad + \overbrace{\theta\int_{p_w=p_l^{\mathrm{IIP}}}^{p_w=p}\left((1-F_w(q))f_l(q)(q-c_l)^{1-r}\right) f^{\mathrm{IIP},\theta}(p_w)dp_w}^{\text{Low-cost rival exists but bidder wins stage 2}} f^{\mathrm{IIP},\theta}(p) \\
&= (1-\theta F^{\mathrm{IIP},\theta}(p))(p-c_l)^{1-r} + \\
&\quad (1-\theta F^{\mathrm{IIP},\theta}(p))(\alpha+(1-\alpha)(1-F_l(q)))\left((p-c_l+q-c_l)^{1-r}-(p-c_l)^{1-r}\right)f_w(q) + \\
&\quad \theta\int_{p_w=p_l^{\mathrm{IIP}}}^{p_w=p}\left((\alpha(c_h-c_l+p_w-c_l)^{1-r}+(1-\alpha)(p_w-c_l)^{1-r})^{\frac{1}{1-r}}-(p_w-c_l)\right)^{1-r} f^{\mathrm{IIP},\theta}(p_w)dp_w
\end{aligned}$$

Simplifying this, we obtain Equation 3. We differentiate that equation with respect to p. The left hand side expression tends to zero since the expectation of the profit from any of the actions in the strategy set is the same. Using that, we obtain a differential equation involving $f^{\mathrm{IIP},\theta}(p_w)$ and $F^{\mathrm{IIP},\theta}(p_w)$, which when solved yields the equilibrium bid distribution. The differential equation does not have a closed form expression and has to be solved numerically.

Appendix B: Risk Aversion Estimates for Dataset 2

		All Treatments	CIP		IIP	
			$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$
Dataset 2	Risk aversion coefficient	0.51**	0.62**	0.26**	0.01**	0.63**
	Stdev	(0.02)	(0.03)	(0.02)	(0.001)	(0.02)
	Observations	2,032	352	664	352	664

Table 7: Results for estimated risk aversion coefficients. Numbers in parentheses indicate the bootstrapped standard errors. Note that ** indicates a significance level of 1%. The first ten periods are omitted from the estimations.

Appendix C: Experiment Instructions (Complete Information Policy, θ=0.9)

This is an experiment in the economics of strategic decision making. Purdue University has provided funds for this research. If you follow the instructions and make appropriate decisions, you can earn an appreciable amount of money. The currency used in the experiment is francs. Your francs will be converted to U.S. Dollars at a rate of _____ francs to one dollar. At the end of today's session, you will be paid in private and in cash.

It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

The experiment consists of 50 decision making periods. Each period you will be grouped with one other person in the experiment. At the beginning of each decision making period you will be randomly re-grouped with another person. Since the groupings change randomly every period, you will be grouped with a new person in almost every period. These instructions are for Part 1, which lasts for 25 periods. You will receive additional instructions for Part 2.

Your Offer Prices and Profits

During each period, you can sell units of a fictitious commodity. If you sell a unit, then you will have to incur that unit's production cost. Each period you and all other participants will make two choices—an offer price in stage 1 and an offer price in stage 2. Each represents an offer price to sell a unit of a fictitious good to the experimenter. You can sell one unit in each of the two stages. If you sell your unit, then you will earn profits (in experimental francs) equal to

Your profits = Your offer price – Your production cost

If you do not sell your unit in a stage, then your profit for that stage is 0. This will happen frequently, since only one of the two people in your group can sell a unit in each stage.

For example, suppose your production cost is 200 and your offer price is 322, and you sell a unit in this stage. Then your profit would be 322 – 200 = 122 for this stage. Note that you only incur your production cost if you sell a unit.

Costs are Determined Randomly

Your costs and the costs of the other person in your group are determined randomly by the computer at the start of each period. Everyone's costs remain unchanged for both stage 1 and stage 2 within a period, but then they are randomly determined again at the start of each period.

There is a 90% chance that your cost is 200 and a 10% chance that your cost is 400. Which cost you have this period is determined through a (virtual) “ball draw” from a bingo cage containing 10 balls, comprised of 9 red and 1 black balls. If a red ball is drawn then your cost is 200 and if a black ball is drawn then your cost is 400. The other person in your group will have a separate ball draw (with replacement) to determine his or her cost. Everyone will have a new ball draw to determine cost at the start of every period. Everyone always simply has a 90% chance (that is, a 0.9 probability) of having a 200 cost. Remember, everyone’s cost also remains unchanged over the 2 stages of each period.

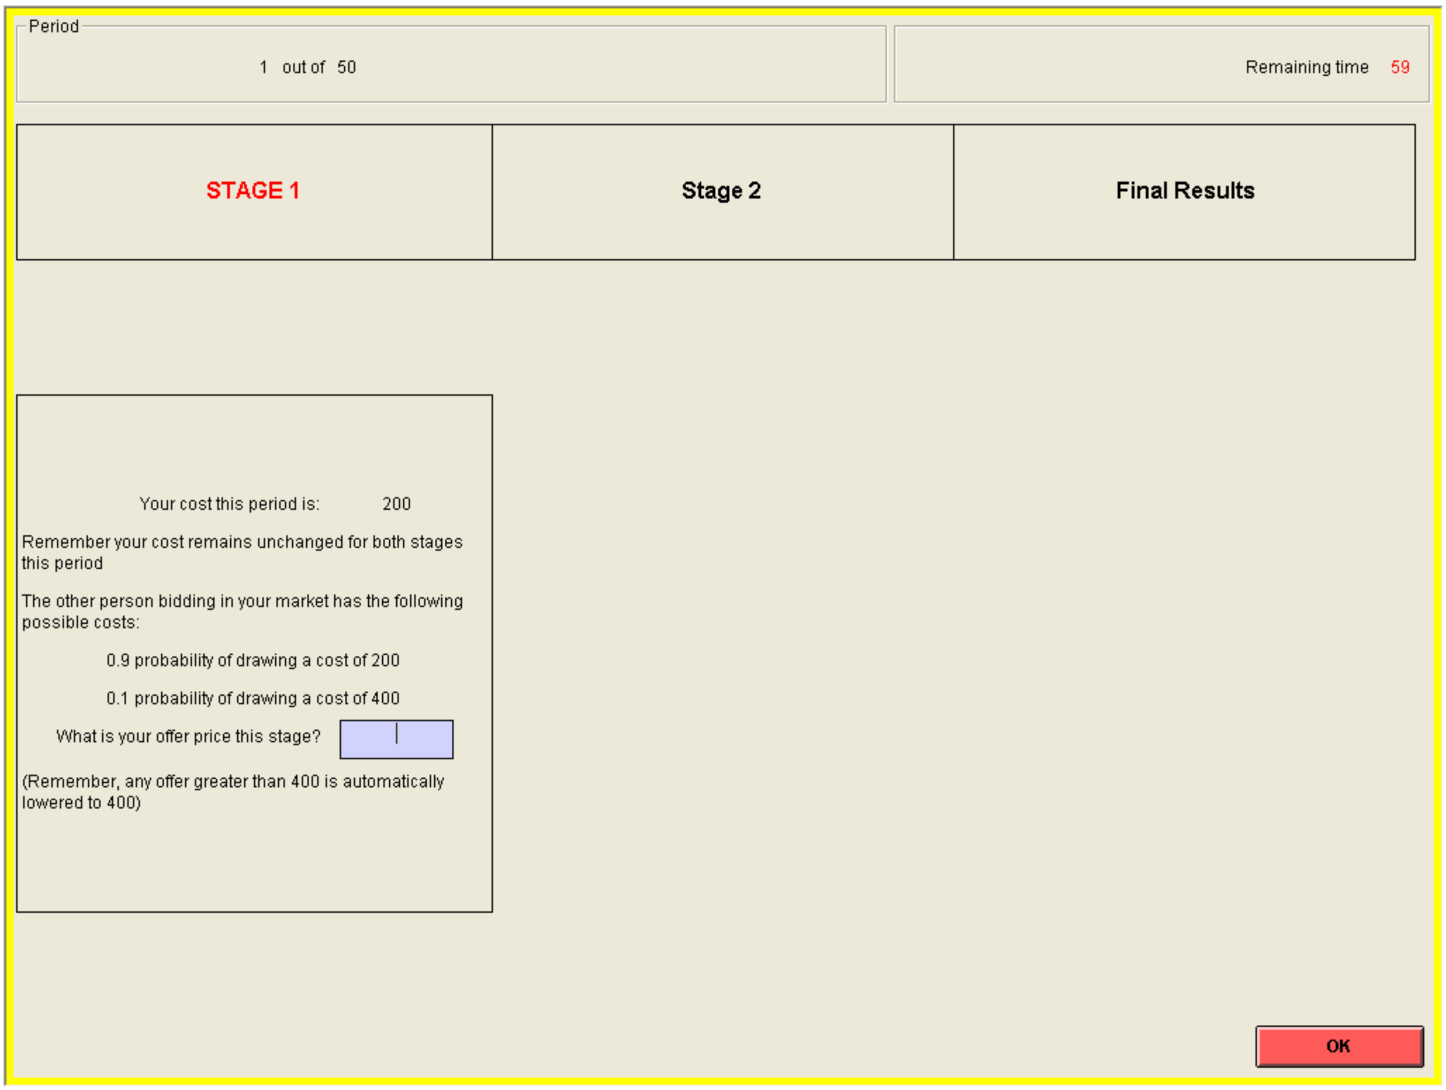

<u>Submitting Your Offer Prices</u>

You will submit your offer prices using your computer. An example screen for Stage 1 is shown above. As you can see on this screen, you will know your cost for the period before you submit your offer, but you will not know the cost (or the offer price) of anyone else at this stage. Up to two decimal places are permitted for any price offer.

Determining Who Makes the Sale

The computer determines whether you or the other person in your group makes the sale each stage following a very simple rule: The lowest offer price in your two-person group sells the unit, as long as this lowest offer price is not greater than 400. Since the computerized buyer will not pay more than 400 for a unit, any offer that you submit that is greater than 400 will be automatically lowered to 400. If the two offer prices are equal, then the person who sells the unit is determined randomly. The buyer will buy at most one unit in each stage from each two-person group.

Stage 2: New Offer Prices

You will learn the Stage 1 offer price submitted by the other person in your group at the start of Stage 2, as shown above. This screen will also indicate who sold a unit in Stage 1. At this point you will submit a Stage 2 offer price, at the same time the other person submits her price. Your cost does not change between the two offer stages.

Guessing the Cost for the Other Person

The other person's cost also does not change between the offer stages. At the same time that you submit your Stage 2 offer price, you will also enter a guess about the chances that the other person has a cost of 200. (Remember, we already told you that costs are determined randomly at the beginning of the period, and everyone always has a 90-percent chance of having the cost of 200 for the period.) What you enter on your screen is the *probability* that the other person has a cost of 200 this period. For example, if you think that she has a 50-percent chance of having a cost of 200, then you enter 0.5. Or, if you think that she is three times as likely to have a cost of 200, rather than the cost of 400, then you enter 0.75. (Up to two decimal places are allowed.) Or, if you think that she certainly does *not* have a cost of 200, then you enter 0.

Your guess can earn you additional money. At the end of the period, we will show you the cost of this other person, and compare it to your guess. We will then pay you for the accuracy of your guess as follows:

Suppose you guess that the person you are grouped with has a cost of 200 with a 75% chance and a cost of 400 with a 25% chance (as in one example above). Suppose further that this person actually has a cost of 400. In that case your

Guess Payoff $= 20 - 10(1\text{-}0.25)^2 - 10(0.75)^2 = 8.75$ francs.

In other words, we will give you a fixed amount of 20 francs from which we will subtract an amount that depends on how inaccurate your guess was. To do this we use the cost of the person you are grouped and we will take the probability you assigned to that cost, in this case 25% on 400, subtract it from 100% and square it. We will then take the probability you assigned to the wrong cost, in this case the 75% you assigned to 200, and square it also. These two squared numbers will then be multiplied by 10 and subtracted from the 20 points that we initially gave you, to determine your final guessing payoff (which is 8.75 francs in this example).

Note that you get the lowest payment under this payoff procedure when you state that you believe that there is a 100% chance that the other person has a particular cost when it turns out that she actually has the other cost. In this case your guessing payoff would be 0, so you can never lose earnings from inaccurate guesses. You get the highest payment if you guess correctly and assign 100% to the cost that turns out to the actual cost of the person you are grouped with; in this case your guessing payoff would be 20 francs.

Note that since your guess is made before you know the cost of the person you are grouped with, you maximize the expected size of your guessing payoff by simply stating your true beliefs about what you think this other person's cost is. Any other guess will decrease the amount you can expect to earn from your guessing payoff.

The End of the Period

After everyone has submitted offer prices for both stages of the current period you will be shown the final results screen, as shown on the next page. This screen displays your offer prices as well as the offer price and cost of the person you are grouped with for the current decision making period. It also shows your total earnings for this period and your cumulative earnings for the experiment so far.

Once the outcome screen is displayed you should record your offer prices, cost, and the other person's offer prices and cost on your Personal Record Sheet. Also record your current and cumulative earnings. Then click on the *continue* button on the lower right of your screen. Remember, at the start of the next period all participants are randomly re-grouped, and you are randomly re-grouped each and every period of the experiment.

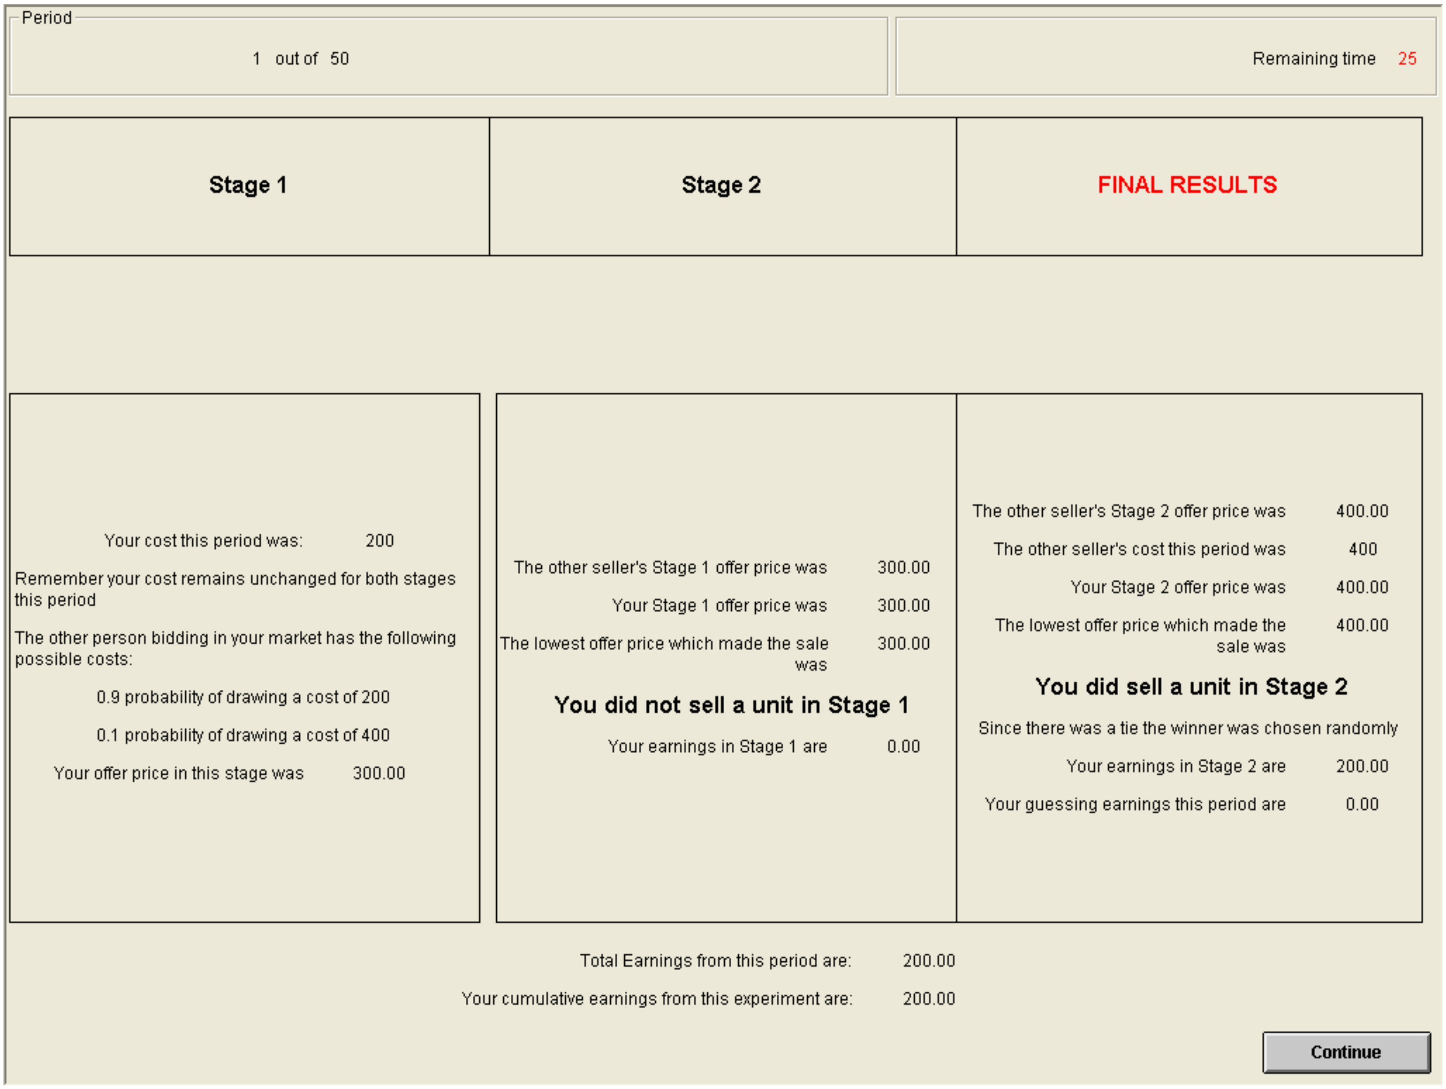

Personal Record Sheet – Part 1

Period	Your Cost this Period	Your Stage 1 Offer Price	Other Seller's Stage 1 Offer Price	Your Stage 2 Offer Price	Other Seller's Stage 2 Offer Price	Other Seller's Cost this Period	Your earnings this period	Total earnings in Part 1 so far
1								
2								
3								
4								
…								
24								
25								