Framing the First-Price Auction

Theodore L. Turocy*

Department of Economics

Texas A&M University

College Station TX 77843

Elizabeth Watson

Department of Economics

Texas A&M University

College Station TX 77843

Raymond C. Battalio†

Department of Economics

Texas A&M University

College Station TX 77843

June 24, 2006

Abstract

We revisit the result that, in the independent private values setting, the first-price sealed bid and descending clock (or Dutch) auction are not isomorphic. We investigate the hypothesis that the empirical non-isomorphism arises from framing and presentation effects. Our design focuses on a careful construction of subject interfaces that present the two environments as similarly as possible. Our sessions also consist of more auction periods to test whether the initial framing effects are subsequently washed away by greater experience. We find the difference between the implementations persists. We also investigate an intermediate implementation which operates like the Dutch auction, but in which the clock continues to tick to the lowest price without informing

*. Corresponding author; email address `turocy@econmail.tamu.edu`

†. Deceased.

bidders when others have bid on the object. This implementation gives results in between the sealed-bid and Dutch implementations.

Keywords: auctions, first-price, isomorphism

---

1 Introduction

Despite being one of the simplest auction environments, the single-unit independent private values first-price auction has exhibited two intriguing stylized facts in laboratory studies. First, subjects typically bid more aggressively than predicted by the risk-neutral Bayes-Nash equilibrium (RNNE for short), resulting in significant reductions in their overall monetary earnings. In addition, the first-price auction may be implemented in two ways, as a sealed-bid or as a Dutch descending-clock auction. Bayes-Nash equilibrium theory suggests these should be isomorphic, giving identical results. Yet, bidding behavior in the Dutch implementation is significantly less aggressive.[1]

The seminal papers in this area include those of Coppinger, Smith, and Titus [3], Cox, Roberson, and Smith [4], and Cox, Smith, and Walker [6], with a substantial subsequent literature. Kagel [9] provides a good survey and summary of the results and the theories put forward to explain them. Instead of testing these theories specifically, we return to these stylized facts from a new perspective, in which more careful consideration of the presentation and framing of these auctions plays a central part. Our

1. We will use the term "first-price" auction to refer generically to the institution, with the term "implementation" referring to one of the ways in which the theoretical structure of the auction is realized.

design considers two aspects of framing in the first-price auction.

We investigate the result that both auction implementations produce market prices in excess of the RNNE by presenting the games in the context of a custom-designed graphical interface. This interface presentsthe information, action, and payoff spaces within a unified rectangular area, visually presenting the interrelations among private values, bids, and earnings. Here, we take advantage of the prevelance of the graphical computer interface today, not available at the time of the classic papers cited above. Our subjects make choices and receive feedback in the same frame, in keeping with principles of interaction design (e.g., Cooper and Reimann [2]). The graphical display is identical across implementations, save the minimal changes necessary due to the rules of the games.[2]

As a further control to keep the environments as similar as possible, and in view of the observations of Katok and Kwasnica [10], we attempt to control for the opportunity cost of subjects' time by choosing the speed of the clock in the Dutch auction in such a way that the typical period and session length for all implementations are approximately the same. We view the results of Katok and Kwasnica as an illustration of the importance of maintaining the dominance of the payoff implications of actions within the session (Smith [12]). In a pilot session not reported in this paper, in which we did not institute this control, our subjects completed 80 sealed-bid auctions in about 30 minutes. In subsequent sealed-bid sessions, a common subject question was whether the session would end sooner, or if more periods would be conducted, if they "bid faster." Thus, we view this control as a significant

2. Specifically, subjects submit bids by clicking on the appropriate price in the sealed implementation, whereas they bid by clicking on a "Purchase" button in the Dutch.

design feature for the purpose of testing isomorphism.

The sessions reported in this paper consist of 60 auction periods, in which subjects participate in the same implementation throughout.[3] This duration is announced to the subjects during the instructions for the session. This design feature has two objectives. First, with a large number of auction periods, the cumulative effects of suboptimal bidding, in terms of foregone earnings, will be more significant. Second, if the behavioral features of these auction environments arise largely from an initial misperception of the tradeoffs made in setting bids, a larger number of periods offers time for subjects to adapt.

Our second framing treatment addresses more directly the difference in the extensive forms of the sealed and Dutch implementations. We contrast the way the uncertainty about other bidders' values and bidding behavior is presented in these two games. Dorsey and Razzolini [7] observe that the choice of a bid in the sealed implementation is similar to the choice of a lottery from a menu, where each lottery i in the menu has some probability p_i of winning some prize q_i, with a prize of zero otherwise, such that the p_i and q_i have an inverse relationship. In the Dutch implementation, those same lotteries are presented in a sequential format. At any point in time, the choice of whether to purchase now, or to allow the auction to continue and the clock to tick down, is similar to the choice between an outcome that is essentially a sure thing – purchase the object now and earn the difference between one's value and the current price with a high probability – versus not purchasing. In the latter case, with high probability, the bidder will face a similar choice at the next price, in which the purchase option has now become more attractive

3. We acknowledge the potential interest in studying transfer across the implementations we study, but choose the across-sessions design to abstract away from those considerations.

because the clock price has fallen. We conjecture that the Dutch presentation helps bidders to recognize the tradeoff between probability of winning and the amount won - the essential tradeoff in the first-price auction - in a different way than the sealed implementation. We will refer to this intermediate implementation as the "silent" implementation.

To investigate this, we introduce a synthetic implementation, intermediate between the sealed and Dutch, in which a clock counts down as in the Dutch implementation, but in which the outcome of the auction is not revealed until the clock reaches the lowest price. Thus, it is functionally similar to the sealed implementation in terms of feedback, in that the results of the auction are not known until all choices are made. At the same time, it shares the property of the Dutch implementation in that it cues the bidder to make assessments at the margin about the choice of bids. The silent implementation allows these to be separated to some degree.

We find that we replicate both of the standard empirical regularities. The subjects bid significantly more than the risk-neutral Nash prediction in both implementations, and subjects in the sealed implementation bid more aggressively than in the Dutch. We also find that both of these results persist over the session. There is a small downward trend in prices over the session, but prices remain well above the risk-neutral prediction even after 60 periods. The differences between the sealed and Dutch implementations remain. The results of this silent implementation fall in between the sealed and Dutch: market prices typically exceed those in the Dutch, but are less than those in the sealed implementation.

The paper is organized as follows. Section 2 describes the design of the experimental sessions. Section 3 outlines the experimental results, both at the market and individual levels. Section 4 concludes with a discussion and future directions.

2 Design

Each experimental session was conducted with 18 subjects recruited from the undergraduate student body at Texas A&M University. The 18 subjects were randomly grouped into two cohorts of 9 subjects each. Within each cohort, each period, subjects were randomly assigned to one of three markets, each consisting of three subjects. The sequence of private values and matching into markets is identical across all sessions, facilitating comparisons across the implementations; the two cohorts differ in those sequences, allowing for some control over sequencing effects. The matching was done anonymously, and no subject ID numbers or other information about which subjects were participating in which markets in which periods was known to the subjects. All interaction among the subjects was mediated via computer in the Economic Research Laboratory at Texas A&M.

Figure 1 shows a screenshot of the subject interface presenting the results of a period. Subjects used the rectangle on the left of the screen to interact with the market. All decisions were made in this area, and feedback from the results of the period was presented in the same area. At the right of the screen is a record sheet, which summarized the subject's history, including their private values and bids, the market prices in the markets in which they participated, and their earnings history. The subjects did not observe the results of other markets in which they did not participate. Since the design investigates the isomorphism between the sealed-bid and Dutch implementations, in order to keep information constant across implementations, no information regarding how others bid was presented, since that is not available in the Dutch implementation. The record sheet by default displayed the results of the last 25 periods, but scroll buttons were available for subjects to review the results from earlier in the session.

In each period, subjects received a resale value for a single unit of a fictional commodity drawn from the range \$0.15 to \$6.00 in increments of \$0.15; therefore, there were 40 possible resale values. The resale values were equally likely and drawn independently across periods and subjects. The environment was pure private values: the subject who purchased the object earned the difference between her resale value and the market price; subjects who did not purchase the unit earned zero for the period. Ties were broken at random.

Bids were constrained to \$0.10 increments, starting at \$0.10; the maximum permitted bid was \$6.20. The bid and value spaces were chosen so that in our environment it is a symmetric Bayes-Nash equilibrium to submit a bid equal to two-thirds of the private value. This is the analog of the unique symmetric Bayes-Nash equilibrium of the first-price auction with risk-neutral bidders when values are distributed uniformly and the permitted bids are the nonnegative real numbers.

We consider three implementations of the first-price auction. In the sealed-bid implementation, subjects observe the realization of their private value, and submit a bid by clicking on the corresponding area of the market rectangle. In the Dutch implementation, a clock price was displayed, which started at

$6.20 and decreased by $0.10 each second; once a subject clicked the button to purchase the object, the clock stopped for that market and the results were displayed to all participants in that market. The silent implementation operates as the Dutch does, except the clock price decreases to zero in every period, with no feedback regarding the outcome of the market is given until that time. Subjects were permitted to submit bids, or click to purchase, above their resale value in all implementations.

The clock speed in the Dutch and silent implementations was chosen with two goals in mind. First, the clock ticked slowly enough that subjects could click to purchase while the clock at their desired price with a high degree of accuracy.[1] Second, the clock speed was such that the overall length of the Dutch sessions would be roughly comparable to the sealed-bid sessions.

The subjects participated in 60 periods, and the length of the session was announced during the instructions. Three sessions of each implementation, sealed, silent, and Dutch, were conducted. No subject participated in more than one session. Some subjects had experience in other sessions run in the laboratory, while others did not, but none had any experience in an auction environment. The instructions were read aloud from a projector screen while subjects followed along on their terminal screens. After the instructions were completed, subjects answered a questionnaire, which was checked for correctness by an experimenter. Incorrect answers on the questionnaire were corrected by pointing the subject to the correct answer in the instructions,

1. We implemented "practice" rounds during the instructions during which subjects were asked to "stop" the clock at a given price. No subjects experienced any difficulty in doing so. Thus we believe, at least in the vast majority of cases, that the subjects were able to register their desire to stop the clock and buy the object at the price they intended.

although very few subjects answered any questions incorrectly on the first try.

3 Results

3.1 Market performance

Result 0. *Prices significantly exceed the risk-neutral prediction in all three implementations.*

There are 1080 market observations for each of the three implementations (60 periods $\times$ 6 groups $\times$ 3 sessions). Table 1 categorizes the observed market prices for each of the implementations depending on whether the observed price exceeded, matched, or was less than the risk-neutral Nash prediction.

For all three implementations, the market prices exceeded the risk-neutral prediction in a significant majority of market periods, which replicates previous results. The realized market price in the sealed implementation exceeded the risk-neutral theory price in fully 96.6% of the market periods. In the Dutch implementation, 86.1% of the market periods resulted in a market price exceeding the risk-neutral prediction. The frequency with which the market price exceeded risk-neutral theory in the silent implementation falls in between at 93.4%.

Implementation	Actual > Theory	Actual = Theory	Actual < Theory
Sealed	1043 (96.6%)	11 (1.0%)	26 (2.4%)
Silent	1009 (93.4%)	26 (2.4%)	45 (4.2%)
Dutch	930 (86.1%)	33 (3.1%)	117 (10.8%)

Table 1. Comparison of observed market prices with those predicted by risk-neutral Nash equilibrium. Numbers listed are instances out of 1080 market observations in which the observed price exceeded, matched, or was less than the risk-neutral prediction.

The null hypothesis that the frequencies in Table 1 come from the same distribution can be rejected. For each market, we construct a multinomial random variable which takes on the values +1, 0, or -1 if the realized market price exceeds, equals, or is less than the theory price, and consider the null hypothesis that the underlying frequencies of the three outcomes are identical across implementations. This null hypothesis can be rejected for each pair of implementations at the .01 level of significance.[1] This suggests that the three implementations can be ordered in terms of the observed prices, which is formalized in the next result.

Result 1. *Prices are higher in the sealed-bid than in the Dutch. Prices in the silent implementation fall in between. The difference in prices is persistent*

1. The chi-squared statistic for comparing the distribution from the sealed and silent implementations is 11.72; for silent versus Dutch, 36.06.

throughout all periods.

The phenomenon of higher prices in the sealed implementation is persistent in this environment, even when subjects are given the opportunity to participate in 60 market periods. Recall that the same sequence of private values and matchings was used in each session. Thus, we can make a direct comparison of market prices in each market j in each period k across the three implementations, for a total of 360 market comparisons. The three panels of Table 2 present this comparison for each pair of implementations, aggregating by which implementation resulted in the larger median price.

Table 2 also presents the same comparisons disaggregated by the first and second halves of the sessions. There is no significant alteration of the rankings across implementations between the halves.

Periods	**Sealed > Dutch**	**Sealed = Dutch**	**Sealed < Dutch**
All periods	267 (74.2%)	41 (11.4%)	52 (14.4%)
Periods 1-30	141 (78.3%)	15 (8.3%)	24 (13.3%)
Periods 31-60	126 (70.0%)	26 (14.4%)	28 (15.6%)

Periods	**Sealed > Silent**	**Sealed = Silent**	**Sealed < Silent**
All periods	189 (52.5%)	65 (18.1%)	106 (29.4%)
Periods 1-30	92 (51.1%)	37 (20.6%)	51 (28.3%)
Periods 31-60	97 (53.9%)	28 (15.6%)	55 (30.6%)

Periods	**Silent > Dutch**	**Silent = Dutch**	**Silent < Dutch**
All periods	225 (62.5%)	54 (15.0%)	81 (22.5%)
Periods 1-30	116 (64.4%)	22 (12.2%)	42 (23.3%)
Periods 31-60	109 (60.6%)	32 (17.8%)	39 (21.7%)

Table 2. Comparison of median prices among the three implementations.

Result 2. *There is evidence that prices decline slightly over the course of the session.*

We ask whether, within each cohort, there is evidence of market prices decreasing over time. We specify a simple model relating the observed price to the risk-neutral theory price and the period number:

$$\text{actual}_{it} = \alpha \times \text{theory}_{it} + \beta \times (\text{theory}_{it} \times t) + \varepsilon_{it}, \tag{1}$$

where t denotes period number, i indexes the market within the cohort, and ε_{it} is a noise term. We estimate this model for each cohort individually, and choose to discard any observations in which the winning bid exceeds the highest value in that market.[2] The coefficient estimates are presented in in Table 3. In 16 of the 18 cohorts, the sign of the point estimate of β is negative, as would be expected if prices tend to fall over time. In the two cohorts for which the point estimate of β is positive, the estimate is not statistically different from zero.[3]

Of the 16 cohorts for which the point estimate is negative, the null hypothesis that $\beta = 0$ can be rejected in 9 cases. In addition, the magnitude of the esti-

2. We are primarily interested in learning at the margin. While learning not to submit bids over value is certainly learning, we want to focus on the extent to which bidding may become less aggressive over time. Furthermore, by throwing out these observations, we only make it more difficult to obtain a negative time trend estimate.

3. The standard errors are corrected for heteroskedasticity using STATA's robust standard errors. Heteroskedasticity is present in the data as there is greater variance in observed market prices in markets where the theory prediction is higher.

mates is small, even when statistically different from zero. Estimates of β are on the order of -0.001, which, over the 60 periods of the session, corresponds to a decrease in the ratio of the observed market price to the theory price of about .06. This still implies significant overbidding relative to the risk-neutral prediction even after 60 periods.

Implementation	**Cohort**	<alpha>$\boldsymbol{\alpha}$	$\boldsymbol{\beta}$	**Adjusted** $\boldsymbol{R^2}$
sealed	20041201.1	1.269	-0.00026 (0.638)	0.991
	20041201.2	1.284	-0.00089 (0.093)	0.992
	20041203.1	1.297	-0.00128 (0.007)	0.995
	20041203.2	1.264	-0.00090 (0.132)	0.990
	20050429.1	1.284	-0.00100 (0.048)	0.993
	20050429.2	1.306	-0.00083 (0.062)	0.996
silent	20050309.1	1.260	-0.00108 (0.067)	0.989
	20050309.2	1.289	-0.00032 (0.565)	0.993
	20050331.1	1.241	-0.00111 (0.039)	0.992
	20050331.2	1.214	-0.00109 (0.068)	0.990
	20050405.1	1.306	-0.00042 (0.421)	0.993
	20050405.2	1.187	-0.00090 (0.111)	0.988
Dutch	20041202.1	1.157	+0.00078 (0.321)	0.986
	20041202.2	1.240	-0.00018 (0.665)	0.994
	20041206.1	1.168	-0.00137 (0.040)	0.980
	20041206.2	1.143	-0.00057 (0.362)	0.987
	20050407.1	1.193	-0.00249 (0.001)	0.979
	20050407.2	1.180	+0.00034 (0.586)	0.990

Table 3. Estimates for the market price model (1). Numbers in parentheses are p-values for the test of the null hypothesis $\beta = 0$ against the two-sided alternative $\beta \neq 0$, with correction for heteroskedasticity.

Result 3. *All implementations obtain high levels of efficiency, and significantly exceed the efficiency predictions of a zero-intelligence random bidding model. The Dutch implementation exhibits slightly lower efficiencies than the other two.*

The efficiency of these markets can be measured in two ways: the percentage of gains from exchange realized, and the frequency with which the highest-value bidder was allocated the object. Tables 4 and 5, respectively, present these measures for each of the three implementations.

To obtain a useful baseline for interpreting these levels, we consider a model of "zero-intelligence" random bidding patterned after that of Gode and Sunder [8] for continuous-time double-auction markets, which was also applied by Cason and Friedman [1] to call markets. In our environment, we operationalize this by assuming bidders choose any individually rational bid with equal probability. Simulation results give a percentage efficiency of 89.4% and a frequency of efficient allocation of 64.3% for this zero-intelligence model. Thus, observed efficiencies do significantly exceed those generated by random bidding behavior.[4]

Tables 4 and 5 also present the evolution of efficiency measures over the course of the 60 periods, aggregated into segments of 15 periods. Significant efficiency improvements are observed early in the session in all implementations, while the magnitude of these efficiency differences across implementations decrease substantially. We also note that while the prices in the silent implementation fall roughly halfway between those in the sealed and Dutch implementations, the efficiency measures of the silent implementation track the sealed imple-

4. See Appendix A for more detail on zero-intelligence efficiency predictions in first-price auctions.

mentation closely.

Implementation	Overall	1-15	16-30	31-45	46-60
Sealed	98.6%	98.0%	99.2%	98.9%	98.6%
Silent	98.6%	98.0%	98.7%	98.9%	98.6%
Dutch	97.7%	96.7%	97.8%	97.9%	98.3%
zero-intelligence	89.5%	89.4%	89.2%	89.6%	89.6%

Table 4. Average percentage of gains from exchange realized, by implementation. The row labeled zero-intelligence presents the expected gains from exchange if all bidders chose bids randomly from their individually rational bids, conditional on the private values used in the sessions.

Implementation	Overall	1-15	16-30	31-45	46-60
Sealed	90.6%	88.5%	92.6%	91.9%	89.6%
Silent	90.4%	89.3%	89.3%	92.6%	90.4%
Dutch	85.7%	79.6%	85.9%	88.9%	88.5%
zero-intelligence	65.7%	63.3%	65.7%	67.5%	65.8%

Table 5. Percentage of markets where highest-value bidder purchased object, by implementation. The row labeled zero-intelligence presents the frequency of efficient allocation if all bidders chose bids randomly from their individually rational bids, conditional on the private values used in the sessions.

3.2 Individual performance

Result 4. *In all implementations, bidders leave significant amounts of earnings "on the table" due to bidding more aggressively than risk-neutral.*

Summary statistics for subject earnings are presented in Table 6, along with the RNNE earnigns predictions.[5]

	Sealed	**Silent**	**Dutch**	**Theory**
Mean	\$12.86	\$14.96	\$17.65	\$30.34
Median	\$12.45	\$15.03	\$17.90	\$29.81
Minimum	\$7.25	\$3.40	\$6.70	\$23.10
Maximum	\$29.25	\$30.00	\$33.15	\$42.10
N	54	54	54	

Table 6. Summary statistics for distribution of subject earnings, by implementation. The theory column refers to the predictions of the risk-neutral Nash equilibrium.

By comparison, we take the risk-neutral Nash equilibrium bidding function as an alternative heuristic which any bidder might have alternatively chosen unilaterally. Figure 2 plots the distribution of earnings foregone by bidders relative to this benchmark. In the sealed implementation, the mean earnings loss over the session relative to this benchmark was \$7.93, with a median of \$8.12. For the silent implementation, the mean and median earnings loss was \$7.38.[6] Thus, most bidders would have been significantly better off unilaterally

5. The earnings totals presented are for the contingent portion of the experiment only, and do not include the \$10.00 participation fee.

using this simple, less aggressive heuristic, even holding constant aggressive bidding by other subjects.

Result 5. *Bidders frequently, though not always, change bids in accordance with a directional learning rule. Directional learning fits the data better in the clock-based implementations.*

Selten and Buchta [11], by using a bid function elicitation approach in a first-price auction, are able to test hypotheses about directional learning. A directional learning rule, applied to the first-price auction, implies that a bidder who wins the object would adjust his bid function at his realized value downward, since he almost certainly could have won the object at a lower price. Similarly, a bidder who does not win the object, but who, in retrospect, could have submitted a bid that would have won the object at positive earnings, would adjust his bid function upwards.

Although we do not directly observe bid functions, there are 26 instances in each session in which a bidder receives the same resale value in two consecutive periods. For these instances, we examine whether the bidder's behavior between those periods is consistent with the directional learning hypothesis. Because independence assumptions across observations in the same session

6. Six subjects overall actually exceeded the benchmark, because ex-post their bidding behavior happened to be superior to the risk-neutral benchmark; of these, four played the role of participant number 7 in their respective sessions, including the significant outlier in the sealed implementation (who earned $6.05, or 61%, more than the risk-neutral benchmark would have earned). This is a feature of the particular realization of values and matchings in the session.

may not be appropriate, we summarize our observations.

Table 7 summarizes the data on how bidders set their bids when receiving the same value as in the previous period. A feature of Selten and Buchta's data is that there was a great deal of inertia: more than 80 percent of the time, their subjects did not adjust their bid function in response to the outcome of the previous period. We observe the opposite: with a similar frequency our subjects do change their bidding behavior.

We conjecture that this result is driven by a bias in the design of Selten and Buchta's interface. In their design, subjects could submit the same bid function as the previous period with a single mouse click, but had to redraw the entire bid function to modify any part of it. In our implementations, the amount of work to set the bid in the subsequent period is the same regardless of whether or not the subject changes their behavior relative to the previous period. Similar to Selten and Buchta, we find that when the subjects do change their bids, they generally do so in the direction that directional learning predicts. Cason and Friedman [1] also investigate directional learning in a double-auction call market, similarly finding stronger evidence for directional learning.

The implementation may have some effect on the frequency with which subjects change their bids. In the silent implementation we observe less inertia and a greater frequency of decreasing bids after having won the previous period. A similar pattern appears in the Dutch implementation. There is a censoring bias in the Dutch auction data, since it is possible to win the auction in one period, and then lose the next period because the market price increases due to another bidder purchasing at a higher price. In these cases, it would be possible for the subject to have intended to increase his bid, but we would not observe this. The row in Table 7 labeled "Dutch (worst case)" presents the worst case for directional learning in this implementation, by assuming that in the 7 cases in which this occurred in the data, the bidder in fact was intending to increase his bid. Even with this worst case scenario, the data generally favor the directional learning hypothesis.

	After winning			**After losing**		
	Increase	**Same**	**Decrease**	**Increase**	**Same**	**Decrease**
Sealed	6 (16.7%)	10 (27.8%)	20 (55.6%)	6 (75.0%)	2 (25.0%)	0 (0.0%)
Silent	4 (11.4%)	5 (14.3%)	26 (74.3%)	9 (56.3%)	5 (31.3%)	2 (12.5%)
Dutch	0 (0.0%)	6 (24.0%)	19 (76.0%)			
Dutch (worst case)	7 (21.9%)	6 (18.8%)	19 (59.4%)			

Table 7. Test of directional learning hypothesis. Numbers in parentheses are percentages.

Result 6. *There is significant heterogeneity in individual bidder adaptation over the course of the session in the sealed and silent implementations, with no overall upward or downward trend either individually or in aggregate. There is weak evidence that bidders exhibit less inertia in the silent implementation than the sealed.*

In parallel with the foregoing directional learning analysis, we consider cases in which the same subject received the same resale value in two periods. Instead of requiring the periods to be immediately subsequent, we require only that the two instances occured within at most 25 periods. Additionally, we require that the resale value was at least \$3.00.[7] Since we are interested in the overall bidding trends, we do not distinguish whether the subject won the auction in the earlier period of the pair.

The data for the 54 subjects in each of the sealed and silent treatments are represented in the simplices in Figure <>3.[8] Graphically, points closer to the apex of the simplex represent bidders who exhibited more inertia, that is, who more often submitted the same bid in both periods of the pair. Points to the left of the vertical line are bidders who, conditional on having changed their bid, increased their bid more often than decreased it. Of the 18 value sequences in the parameter set, the number of period pairs satisfying our restriction ranges from 6 to 14, with a mean of 10.8 and a median of 11.5.

In the sealed implementaiotn, we observe 23 bidders who increased their bid in the second period of the pair more often than decreased it, and 26 bidders who decreased more often than increased, with 5 bidders equally likely to

7. The cutoff of 25 periods is chosen because the subjects were shown the results of the last 25 periods (including the current period) on their record sheet without having to scroll. While we did track when and how often subjects made use of the scroll buttons within the software, we suspect scrolling was done just as much to fill idle moments as to garner information, and so we do not attempt any serious interpretation of that data. We choose \$3.00 as a cutoff for the private value to consider only cases in which the subject could reasonably be said to expect to be competitive in the auction. Adjusting or removing these cutoffs do not change the qualitative results.

8. We omit analysis of the Dutch implementation for the same bias reasons as arose in the directional learning analysis. Of the 585 possible period pairs matching our restriction in the parameter set, in only 268 instances did the bidder win the object in both period. Thus, the number of censored observations is significant, rendering any useful conclusions from this data impossible.

go either way conditional on making a change. In the silent treatment, 22 bidders increased more often than decreased, and 26 decreased more often than increased, with 5 bidders equally likely to adjust in either direction.

Graphically, the points cluster lower in the simplex for the silent implementation. This indicates that bidders overall tend to change their bids more in the silent implementation than in the sealed implementation. This observation can be made more precise by considering the empirical cumulative distribution functions of the frequency with which bidders in each implementation kept the same bid. The distribution of this frequency across bidders in the sealed implementation first-order stochastically dominates that of the silent implementation. In terms of the representation in Figure 3, this is equivalent to saying that for any line drawn horizontally across the two simplices at the same height, there are more bidders whose points fall on or below that line in the silent implementation graph than in the sealed implementation graph.

4 Discussion

We replicate two empirical regularities in laboratory implementations of first-price auctions with private values. A sealed-bid implementation generates higher market prices than a Dutch implementation, even though the two are considered isomorphic in theory. Furthermore, both implementations give market prices in excess of those predicted by risk-neutral Bayes-Nash equilibrium.

Our design extends these replications in some new directions. These differences are persistent over a session of 60 market periods. Thus, advance knowledge of the relatively large number of periods does not seem to encourage subjects to recognize that bidding less aggressively is in fact (expected) earnings-enhancing, nor does repetition of the auctions lead to this realization. In addition, an interface in which the magnitude of earnings is prominently displayed in the feedback space does not encourage subjects to bid less aggressively.

To investigate the failure of isomorphism, we introduce an intermediate synthetic implementation in which the market operates like the Dutch auction with a descending clock, but in which the choices of the bidders are not revealed until the clock reaches the lowest price, simulating a sealed bid. We find that market prices in this implementation generally lie between those generated by the Dutch and the sealed bid. We interpret this to suggest that the way the Dutch auction presents the tradeoff between the probability of winning the auction and the amount a bidder pays does help focus at least some subjects on considering this tradeoff more explicitly. However, the silent implementation can be viewed as presenting this tradeoff as a hypothetical, since it is always possible that some other bidder has already submitted the winning bid; we cannot be sure how subjects' behavior is changed due to this hypothetical nature of the presentation. Thus, while we use the sealed implementation as a sort of proxy for determining how individual behavior changes in a clock-based implementation of the first-price auction, we can only conjecture that this is how subjects would also adapt in the Dutch were we able to observe their censored, intended behavior.

Subject to this caveat, we note that the clock-based implementations appear to encourage subjects to be more active in adapting their behavior over time.

However, no clear patterns emerge from this observation: most subjects are just about as likely to increase or decrease bids, even though in almost all cases their initial bids are too high for an expected-money maximizer.

To the extent that one purpose of laboratory economics is as a testbed for mechanisms that will subsequently be deployed in the field, one would like to know what conditions are necessary to recover a reasonable approximation of a theory's predictions. In the case of Bayes-Nash equilibrium, the theory is silent on the details of the presentation of the game, insofar as the bidders are expected to reduce the compound lottery presentation of the Dutch auction in the same way as the menu lottery presentation of the sealed-bid. While on the one hand it is attractive to think of the Dutch auction as being more explicit in presenting to the subjects the key marginal price-versus-probability tradeoff central to the formulation of bids, it would be a stretch to conclude the Dutch implementation might be somehow eliciting bidding behavior from the subjects that more closely approximates "optimal" behavior given their preferences.

Finally, we conclude with a few casual observations. When we distributed screenshots of the instructions to the subjects, we instructed them that they were "free to mark up those pages in whatever way they might find helpful." While most subjects made no marks or doodled, a few chose to attempt certain forms of market analysis. A few subjects chose to track their earnings per minute, which we interpreted as further evidence that controlling the length of the session is a significant aspect of testing the isomorphism. No subjects tracked the most relevant datum - the distribution of the market prices they observed - though several attempted to approximate the average market price.

However, among those subjects, none attempted to distinguish periods in which they purchased the object from those in which they did not. Note that those market statistics, then, are generated by the maximum bid out of *three* bids, whereas in computing an optimal bid in a three-bidder market, the subject is interested in the assessed distribution of the maximum bid out of the *two other* bids in the market. While we do not assert a direct link between these statistics and the way those bidders formulated their bids, we note that failure to correct for one's own effect on the history of market prices would also lead to more aggressive bidding.

Dorsey and Razzolini [7] investigate bidding behavior against robots programmed to bid according to the risk-neutral Nash equilibrium in four bidder sealed-bid auctions. They consider treatments in which the interface does compute the probability of winning with any given bid, and those in which it does not. For high realizations of the private value, the range for which the choice of bid is most important, they find that giving the subjects the probability of winning makes bidding less aggressive. Our hypothesis is consistent with their results in that when subjects generate impressions of the pattern of high bids, they fail to correct for the fact they themselves are part of the process that generates those high bids.

Appendix A Zero-intelligence efficiency

The ZI efficiency predictions we present for our environment are not significantly affected by our discretization of the auction environment. For the environment used by Cox et al [4], the ZI model predicts a percentage efficiency of 89.0% and a frequency of efficiency allocation of 64.1%. For continuous independent private values, the predictions of the ZI model are 89.3% and 62.9%, respectively. So, these figures are in the same ballpark for most environments which have appeared in the literature.

Changing the number of bids and values in a discretization has opposite effects on the two efficiency measures under the ZI model. On the one hand, percentage efficiency increases as one approaches the continuous limit. To see why this happens, take two values v_1 and v_2 which are adjacent in a discretization, and consider a market in which the two highest realized values are v_1 and v_2. If an inefficient allocation occurs, in the ZI model the bidder with v_2 is most likely to be purchasing the object. Now, consider a refinement of the discretization with another value v_3 between v_1 and v_2. In addition to keeping all the private value realizations in the original discretization, we have dded realizations in which v_1 and v_3 are the highest realizations. Again, the most likely misallocation in this market under ZI is to the bidder with value v_3. Percentage-wise, this is more efficient than the misallocation to the bidder with v_2 in the original csae.

In contrast, frequency of efficient allocation increases as the number of bids and values possible is reduced. In the continuous limit, the probability the highest value is shared by two or more bidders is zero. As the number of values is decreased, the probability of a realization in which the highest value is shared by two or more bidders increases; in the limit of only one value, the probability of this occurring goes to one. In such cases, allocating to any of those bidders with the highest value is efficient. Now a lower-valued bidder must beat two (or more) higher-valued bidders in order for inefficient allocation to occur; under ZI, this is even less likely to cocur than if the lower-valued bidder needs only to beat one higher-valued bidder.

Acknowledgements

Ray Battalio passed away as the first of the reported sessions were run. Ted and Beth are grateful for his guidance in formulating the research program this paper kicks off.

The instructions, screenshots, and data from this paper are available on the web at `http://econweb.tamu.edu/turocy/papers/isoauction.html`, or by request.

The authors thank John Van Huyck and the Economic Research Laboratory at Texas A&M University for use of facilities and logistical support in running the sessions; Jason Aimone and Chad Wade for assistance during the sessions; Stefan Jacewitz for replicating and verifying the quantitative results; and participants at the 2005 International Meetings of the Economic Science Association for many helpful comments.

References

[1] Timothy N. Cason and Daniel Friedman. Price formation in single call markets. *Econometrica*, 65:311–345, 1997.

[2] Alan Cooper and Robert Reimann. *About Face 2.0: The Essentials of Interactions Design*. Wiley, Indianapolis, 2003.

[3] Vicki M. Coppinger, Vernon L. Smith, and Jon A. Titus. Incentives and Behavior in English, Dutch, and Sealed-Bid Auctions. *Economic Inquiry*, 18:1–22, 1980.

[4] James C. Cox, Bruce Roberson, and Vernon L. Smith. Theory and behavior of single object auctions. In Vernon L. Smith, editor, *Research in Experimental Economics*, volume 2, pages 1–43. JAI Press, Greenwich CT, 1982.

[5] James C. Cox, Vernon L. Smith, and James M. Walker. A test that discriminates between two models of the Dutch-first auction non-isomorphism. *Journal of Economic Behavior and Organization*, 4:205–219, 1983.

[6] James C. Cox, Vernon L. Smith, and James M. Walker. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty*, 1(1):61–99, 1988.

[7] Robert Dorsey and Laura Razzolini. Explaining Overbidding in First Price Auctions Using Controlled Lotteries. *Experimental Economics*, 6:123–140, 2003.

[8] D. Gode and S. Sunder. Allocative efficiency of markets with zero intelligence (ZI) traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101:119–137, 1993.

[9] J. H. Kagel. Auctions: A Survey of Experimental Research. In J. H. Kagel and A. E. Roth, editors, *The Handbook of Experimental Economics*, pages 502–585. Princeton University Press, Princeton NJ, 1995.

[10] E. Katok and T. Kwasnica. Time is money: The effect of clock speed on seller's revenue in Dutch auctions. Working paper, 2003.

[11] R. Selten and J. Buchta. Experimental sealed bid first price auctions with directly observed bid functions. In D. Budescu, I. Erev, and R. Zwick, editors, *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Erlaaum Ass., 1998.

[12] Vernon L. Smith. Microeconomic systems as an experimental science. *American Economic Review*, 72:923–955, December 1982.

[13] Theodore L. Turocy. *Computation and Robustness in Sealed-Bid Auctions*. PhD thesis, Northwestern University, Evanston IL, 2001.

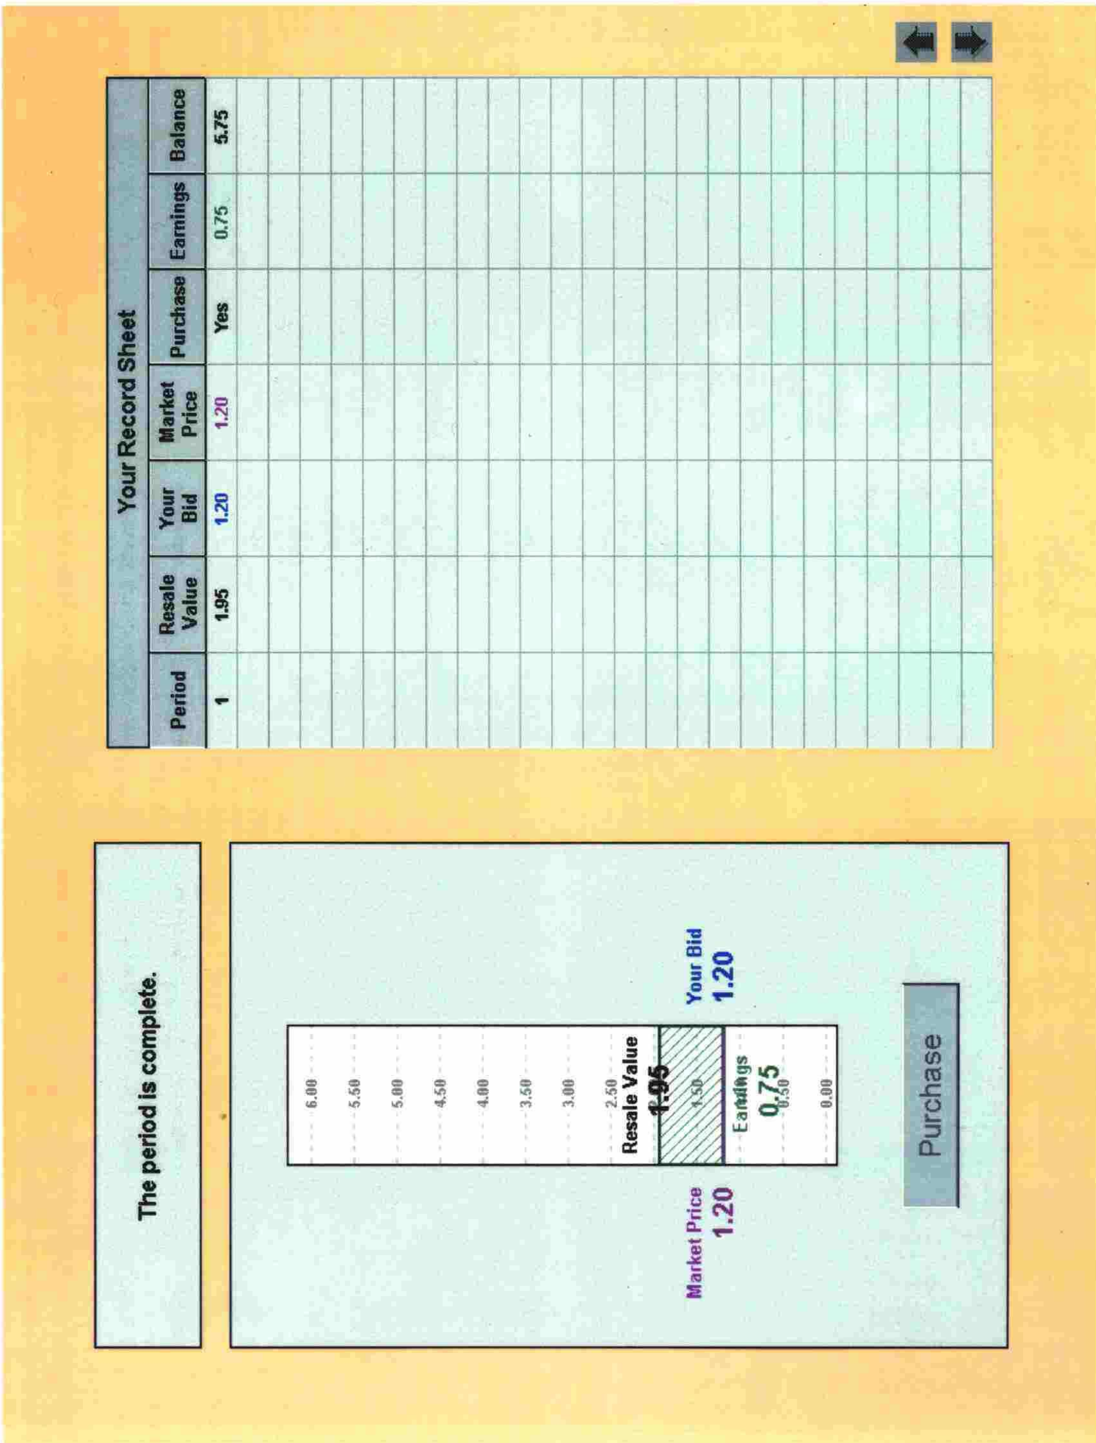

Fig. 1. Screenshot of the subject interface, displaying the result of a period.

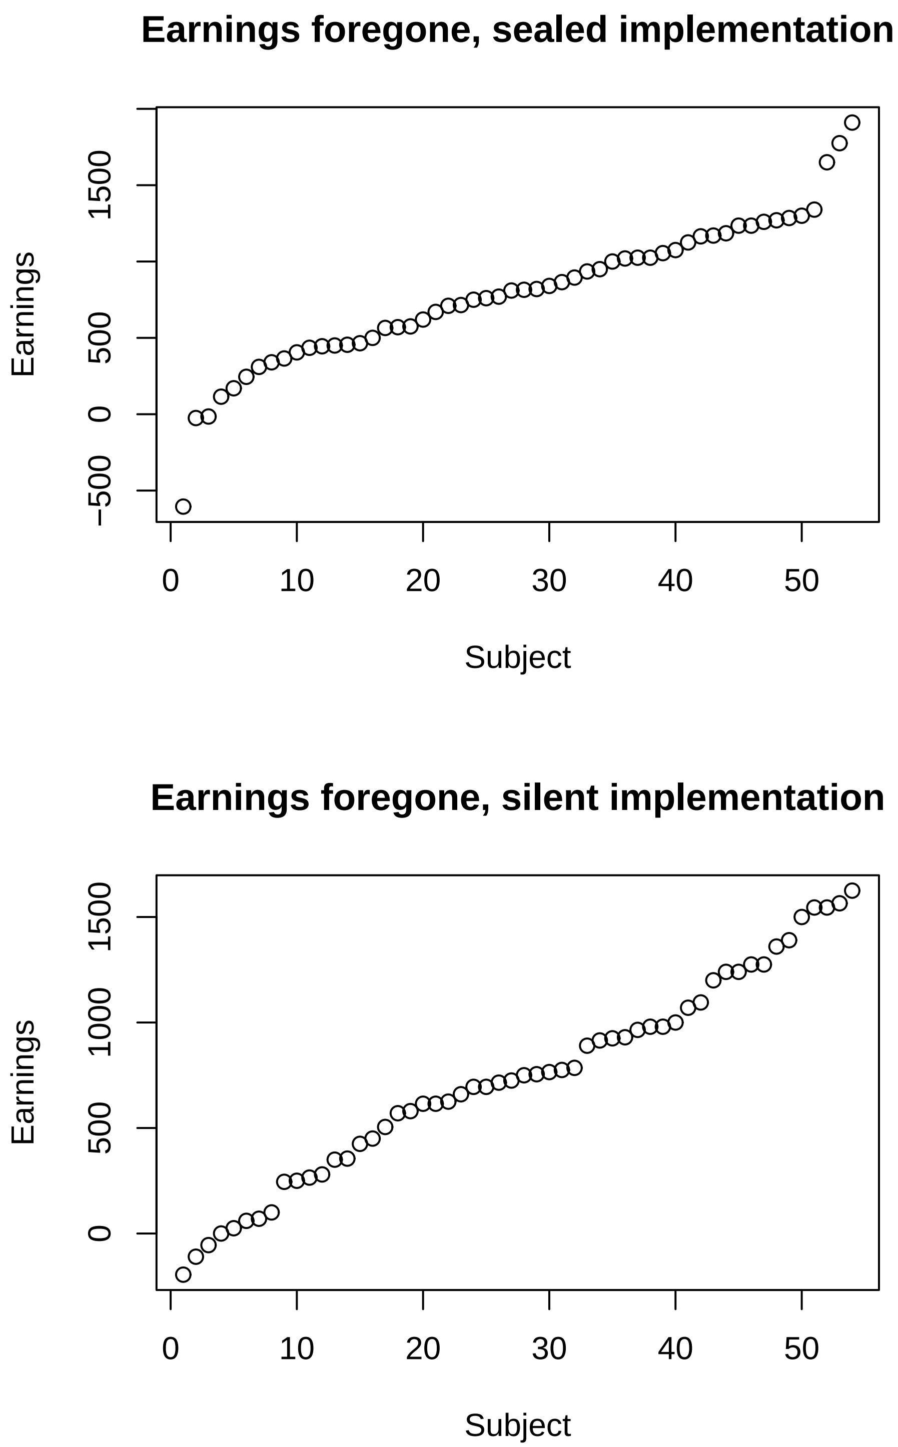

Fig. 2. Earnings foregone by subjects in the sealed and silent implementations, relative to the ex-ante baseline of bidding the risk-neutral equilibrium strategy. Subjects are sorted in increasing order of earnings foregone.

Sealed implementation

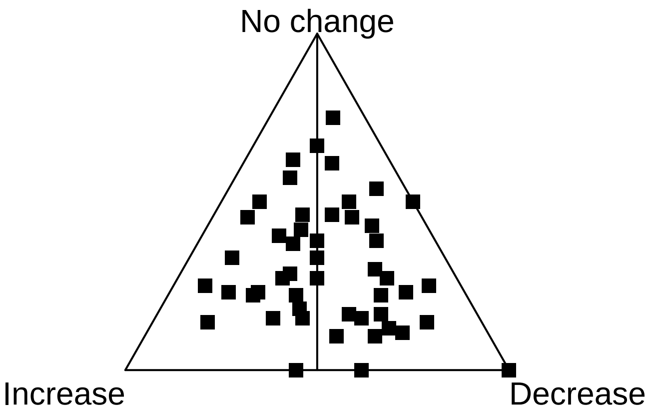

Silent implementation

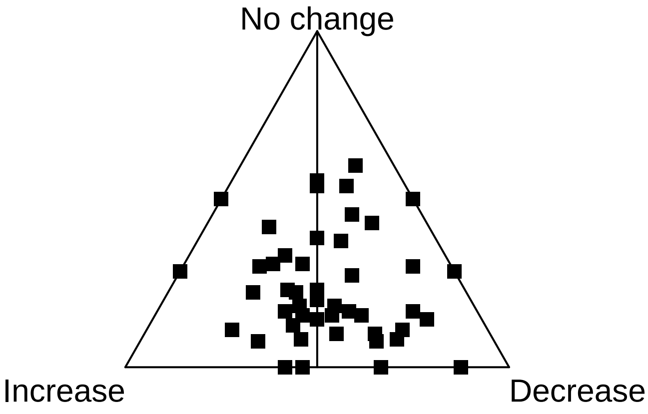

Fig. 3. Frequency with which subjects increase, decrease, or do not change bids, when receiving the same value a second time within 25 periods. Each point represents one subject.