Experimental Economics, 3:261–280 (2000)

Dominant Strategy Adoption and Bidders' Experience with Pricing Rules

RONALD M. HARSTAD
Faculty of Management and RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway, NJ 08854-8003, USA
email: harstad@rutcor.rutgers.edu

Abstract

Truthful revelation is a dominant strategy in both English (oral ascending bid) and second-price sealed-bid private-values auctions. Controlled observations of English auctions are largely consistent with the dominant strategy prediction, but laboratory second-price auctions exhibit substantial and persistent overbidding, even with prior experience. However, the experience of having bid in an English auction has a significant learning effect, moving bidding in subsequent second-price auctions closer to the dominant strategy. I explore two treatments isolating facets of the lessons learned from English auction participation, leading me to the following conclusions. Part of the lesson carried over appears to be considering prices one-by-one, but most of it appears to be a crude awareness that overbidding leads to losses. A claim that English auction experience teaches subjects to recognize the dominant strategy in second-price auctions seems overly optimistic. I introduce a nonparametric technique to test coefficient restrictions when the assumption of normally distributed errors is untenable.

Keywords: auctions, experiments, dominant strategy, nonparametric coefficient, restriction test

JEL Classification: D44, C9

1. Introduction

I report on persistent irrational behavior in a simple experiment by subjects given sufficient financial motivation, and on treatments external to the experimental market that ameliorate this behavior, in ways that feedback does not. The behavior is overbidding, bidding higher than the known value of an abstract asset, in sealed-bid auctions conducted under second-price rules. This fundamentally violates both fully rational and boundedly rational models of behavior, since the observed overbidding is a failure to adopt a dominant strategy.

Let me quickly define terms. The environment I study throughout this paper is *private-values*: a single abstract asset is auctioned each period, with each bidder knowing for sure what the asset is worth to him, and only uncertain of its value to his rival bidders. The game-theoretic prediction is a *dominant strategy*: the predicted strategy yields a higher payoff than any other strategy no matter what strategies govern rivals' bidding.[1] An *English auction* is a dynamic auction in which prices are called out ascendingly, with the last remaining bidder acquiring the asset at the price where his last rival ceased competing. In an English auction, the strategy of dropping out when the price reaches a bidder's value is a dominant strategy (given private values). Vickrey (1961) formalized a *second-price auction* as the sealed-bid equivalent of an English auction: bidders simultaneously and privately submit sealed bids, knowing the rule that the highest bidder acquires the asset at a price equal to the second-highest bid. It is a dominant strategy to bid your value in a second-price auction; this is called *truthful revelation.*

Adoption of the dominant strategy by all subjects is prompt in English auctions, but substantial overbidding by nearly all subjects in second-price auctions persists. We first reported this observation in Kagel, Harstad, and Levin (1987); this sequel attempts some behavioral explanation of the observation and its impact.

My hypothesis to explain this difference in dominant strategy adoption is that the two institutions create differences in feedback about which game-theoretic models are silent. When a subject loses money because both their bid and the second-highest bid exceed their value, this is negative feedback suggesting they may be bidding too aggressively. However, in a second-price auction, a subject might overbid, win, and still make money: it may happen that no rival bids between his overbid and his value. Such an occurrence may be viewed (mistakenly) as positive feedback, serving to offset the negative feedback if overbids are not too large. No corresponding occurrence is possible in an English auction, as the question of whether to overbid does not really arise until the price has reached a subject's value and competitors remain.

It is important to understand this difference in observed behavior for at least four reasons. One is the descriptive relevance of basic game-theoretic notions of behavior, of which adoption of dominant strategies is surely among the most primitive. Another is the tendency of theorists to motivate an analysis of some situation in terms of an English auction, and then choose to conduct the analysis in terms of a second-price auction (dating at least to Milgrom, 1987, called "standard" by Krishna and Rosenthal, 1996). A third reason is to begin to understand what makes a dominant strategy transparent, as a subject is likely to adopt a dominant strategy if the strategy and its dominance *occur* to him.[2] Fourth, it would be useful to establish the limits of truthful-revelation mechanisms in the laboratory.

Several experiments intended to elicit home-grown valuations have used second-price auctions as if their theoretical property of truthful revelation were behaviorally reliable (Hoffman et al., 1993; Shogren et al., 1994, and others); ongoing litigation over the costs of oil spills is but one example why revelation techniques are important. Evidence here offers clear guidance to those researchers as to what sort of auction mechanisms to use, and/or what sort of experience to give subjects in training sessions, in order to retain the possibility that values are being truthfully elicited.[3]

This paper reports laboratory treatments that make notable strides toward an understanding of second-price auction behavior and of the limits to the predictability of dominant strategies. Perhaps most striking is the finding that subjects with experience in English auctions appear to transfer much of what they learned in that institution to later second-price auctions. Further exploration of the same mechanism, experience in another type of auction before bidding in second-price auctions, suggests that some of the lesson transferred to second-price auctions may be the tendency to consider whether individual prices are acceptable, rather than considering how much to bid as a single decision. Specifically, subjects are first given experience in an artificial dominant-strategy mechanism I call a Price Acceptance

List auction, which has no English auction dynamics, then bid in second-price auctions. I also present experiments in which subjects initially have experience bidding in first-price auctions. These indicate that the lesson transferred may merely be that overbidding yields losses, a much cruder notion than that of a dominant strategy.

As is common in experiments, the data defy an assumption of normally distributed errors. Accordingly, my statistical analysis (Section 4) appears to be original in relying on nonparametric tests of coefficient restrictions on estimated bid functions.

2. The laboratory economies

In all experiments, private values for an abstract asset were induced upon subjects, each round, via the following two-step procedure. First, the server computer drew a parameter C from a uniform distribution with support $[L, H]$; the distribution was known to subjects, but not the draw. Then each subject i privately observed a value v_i, an independent draw uniformly from $[C - R, C + R]$, explained as a "resale value." All distributional information, including R, was publicly announced. Note that a subject never faced any uncertainty about the auctioned asset's value to him. This system leaves resale values correlated with each other (technically, affiliated), which has the advantage that a subject almost never learns that his resale value is likely to be among the lower values; knowing this tends to produce "throwaway bids" (unmotivated behavior which can interfere with data analysis) when values are independently drawn.

Subjects privately submitted bids. When all bids were in, all were recorded and posted, and the high bidder's profit was publicly computed, in accordance with the pre-announced pricing rule. All other bidders earned zero profit that period, and the experiment continued to another period with new, independent draws of C and then the v's. Parameters are summarized in Table 1.[4]

Subjects were endowed with a \$5.00 starting capital balance, with profits added to and losses subtracted from it, and the final balances paid in cash. If the current capital balance ever fell below zero, that subject was no longer allowed to bid. Most experiments discussed began with one or two more subjects present than the number bidding, and with subjects taking turns sitting out an auction period, so that the number of bidders could be maintained in the event of a bankruptcy (in Table 1, S is the number of subjects starting the experiment, N is the number of active bidders).

The first five experiments observed inexperienced subjects bidding in second-price auctions. Experiment #6 brought back bidders from #1 and #2, to observe bidders who had an extensive amount of experience in second-price auctions, but no experience in the other auction rules.

English auctions discussed were conducted by calling out prices ascendingly, until all ceased competing but one bidder had dropped. In experiments #7 and #8, active bidders were indicated by holding up signs. Experiment #9 maintained anonymity: the number of active bidders was continually announced, but their identity was kept private, with a system of lights informing only the experimenters as to which bidders were active.

Experiments #10–#16 were designed to observe the effect upon second-price bidding of a small amount of experience (11–16 periods) in another type of auction. Experiments

Table 1. Experimental procedures.

ID#	S	N	Type	R	(Periods)	Remarks
1	8	6	2nd-P	6	(3–8)	Inexperienced subjects
				12	(9–17, 25–29)	Reported bids,profit
				24	(18–24, 30–32)	
2	7	6 (1–23)	2nd-P	6	(3–8)	Inexperienced subjects
		5 (24–33)		12	(9–18, 27–30)	Reported bids,profit
				24	(19–26, 31–33)	
3	7	6 (1–6)	2nd-P	6	(3–8)	Inexperienced subjects
		5 (7–12)		12	(9–16)	
		4 (13–16)				
4	6	6 (1–2)	2nd-P	6	(2–8)	Inexperienced subjects
		5 (3–14)		12	(9–27)	Reported bids,profit
		4 (15–26)		24	(28–32)	
		3 (27–32)				
5	6	6 (1–2)	2nd-P	12	(2–19)	Inexperienced subjects
		5 (3–10)		16	(20–26)	Independent values
		4 (11–30)		24	(27–30)	
6	6	6 (1–13)	2nd-P	12	(2–9, 26–34)	Experienced in 2nd-P
		5 (14–34)		24	(10–17)	Re-recruited from #1, #2
				36	(18–25)	
7	8	6	Engl.	6	(2–6)	Inexperienced subjects
				12	(7–11)	
8	6	6	Engl.	6	(2–7)	Inexperienced subjects
				12	(8–14)	
9	7	6	Engl.	6	(2–6)	Inexperienced subjects
				12	(7–13)	Anonymity preserved
				24	(14–16)	
10	8	6	2nd-P	12	(2–3)	Experienced in English
						Immediately follows #7
11	6	6	2nd-P	12	(2–4)	Experienced in English
				18	(5–10)	Immediately follows #8
12	7	6	2nd-P	6	(2–7)	Experienced in English
				12	(8–15, 24–27)	#9 subjects, next day
				24	(16–23, 28–32)	
13	6	6 (1–10)	P-List	6	(2–8)	Inexperienced subjects
		5 (11–16)		12	(9–16)	
		5 (17–34)	2nd-P	12	(17–27)	Experienced in P-List
				24	(28–34)	
14	6	6 (1–9)	P-List	6	(2–8)	Inexperienced subjects
		5 (10–15)		12	(9–15)	
		5 (16–21)	2nd-P	12	(16–27)	Experienced in P-List
		4 (22–30)		24	(28–30)	
15	8	6	1st-P		(1–14)	
			2nd-P	12	(16–27)	Experienced in 1st-P
				24	(28–34)	
			1st-P		(35–46)	

(*Continued on next page*.)

Table 1. (*Continued*).

ID#	S	N	Type	R	(Periods)	Remarks
16	6	6	1st-P		(1–16)	
			2nd-P	12	(18–28)	Experienced in 1st-P
				24	(29–35)	
			1st-P		(36–43)	

#10–#12 employ subjects who have bid in English auctions, immediately prior or the day before. Experiments #13–#16 try to sort out elements of the experience which subjects may have obtained by participating in English auctions.

For experiments #13 and #14, I created an auction form intended to be more like an English auction, but without any dynamic aspects. I call these "Price Acceptance List" auctions (abbreviated P-List). Each period, each subject learns his value. Each subject is handed a list of 101 prices (uninformatively expanding on the range of conceivable market prices), e.g., \$90, \$89.50, \$89, . . . , \$40.50, \$40. Simultaneously and privately, each subject indicates whether each of the 101 prices is a price that is acceptable for purchasing the asset ("yes" or "no"). Subjects could mark two prices with a "Y" (resp., an "N"), and then draw a line between them to indicate that all prices in that range were acceptable (resp., not acceptable). Note that no suggestion was made as to the order in which subjects might consider various prices. These lists were then submitted, and were in essence used to script an English auction (which was never played out, as that would have suggested a lower-prices-first dynamic). That is, the bidder who labeled a higher price acceptable than any rival obtained the asset, and paid a price equal to the highest price any rival marked yes.[5] After 15–16 periods of experience with P-List auctions, the subjects were switched to bidding in second-price auctions.

The behavioral difference between the English and second-price auctions may relate to the dynamic nature of the English auction, with bidders dropping out as prices rise. Alternatively, it may relate to the English auction focusing attention not on how much to bid, but on whether each particular price is acceptable. The P-List auction carries only the second possibility, and experiments #13 and #14 help isolate the differential impact of this experience.

One specific aspect of the dynamic nature of English auctions may be that subjects seeing the price rise above their value are shown in a particularly transparent way that bidding above their value can only lead to losses. Experiments #15 and #16 explore the role of this crude lesson, in a way unrelated to the fundamental logic of dominance. These experiments observed subjects with a small amount of experience in first-price auctions, where an overbid if highest necessarily loses money. However, first-rice auctions do not have a dominant strategy. After 14–16 periods, subjects were switched to about 20 periods of bidding in second-price auctions, then switched back to first-price auctions.

3. Observations

Table 2 summarizes observations, relative to the dominant strategy, across the classes of bidder experience. Table 3 disaggregates by experiment, and includes data excluded from

Table 2. Summary: Overbidding, by type of experience (expressed as a fraction of equilibrium profit).

Mean	Std. Err.	Std. Dev.	R.P.S.E.	No. Obs.	No. Subjs.
Second-price auctions, inexperienced subjects					
0.415	0.119	2.711	2.742	520	34
English auctions					
0.013	0.067	1.012	1.012	228	21
Price acceptance list auctions					
−0.155	0.389	4.939	4.936	162	12
Second-price auctions, with experience in:					
— Second-price auctions					
0.479	0.050	0.671	0.825	177	6
— English auctions					
0.211	0.053	0.841	0.868	252	21
— Price acceptance list auctions					
0.308	0.050	0.546	0.555	120	12
— First-price auctions					
−0.023	0.101	1.498	1.498	222	14

Table 2 on 4-bidder auctions, which may be less comparable.[6] Throughout, I pooled observations across different numbers of bidders (N) and across different values of R by using the following normalization: $x = (b - v)/[2R/(N + 1)]$.[7] Here, x is the normalized overbid, b the observed bid, v the bidder's value (thus, the predicted bid). The bracketed denominator is the normalization factor, which equals predicted profit to the winning bidder, or, equivalently, the mean difference between the highest and second-highest values. This normalization helpfully focuses on the feedback subjects get: an overbid which is less than 1.0 (normalized) will still yield nonnegative expected profit; greater overbidding creates losses on average.[8]

In Tables 2 and 3, the first column lists the mean (normalized) overbid.[9] Following that, to the right are listed the standard error of the mean and the sample standard deviation. The number shown as R.P.S.E. is the Root Prediction Squared Error, calculated like a standard deviation, except with the prediction (0.0) substituted for the sample mean. Thus, while the standard deviation is a measure of sample heterogeneity, the R.P.S.E. is a measure of the variability around the prediction. The numbers of observations in the sample (and number of subjects in the treatment) are also shown.

A. Second-price auctions. Figures 1–3 show the time path of overbidding across periods for experiments #1, #2, and #6, the latter re-recruiting subjects from the first two. As such, these figures illustrate the temporal path of bidding in second-price auctions as experience is acquired. The figures are identically scaled for ease of comparison. Each period, the mean (normalized) overbid is shown as a "+"; dashes above and below give an indication of the ongoing variability behind these means.[10]

Table 3. Overbidding by experiment (expressed as a fraction of equilibrium profit).

Exp. #	Mean	Std. Err.	Std. Dev.	R.P.S.E.	No. Obs.
Second-price auctions, inexperienced subjects					
1	0.397	0.067	0.906	0.989	180
2	−0.309	0.253	3.338	3.352	174
3	1.175	0.201	1.678	2.048	70
4	1.145	0.325	3.452	3.637	114
5	0.294	0.059	0.665	0.727	126
English auctions					
7	0.036	0.144	1.112	1.113	60
8	0.040	0.134	1.185	1.185	78
9	−0.025	0.078	0.743	0.743	90
Price acceptance list auctions					
13	0.139	0.180	1.639	1.645	84
14	−0.472	0.786	6.894	6.910	78
Second-price auctions, with experience in:					
—Second-price auctions					
6	0.479	0.050	0.671	0.868	177
—English auctions					
10	0.001	0.073	0.253	0.253	12
11	0.170	0.127	0.935	0.950	54
12	0.237	0.061	0.835	0.868	186
—Price acceptance list auctions					
13	0.204	0.053	0.497	0.537	90
14	0.428	0.054	0.439	0.613	66
—First-price auctions					
15	0.011	0.053	0.569	0.569	114
16	−0.061	0.199	2.066	2.067	108

The figures show a clear pattern: subjects upon encountering a second-price auction respond with tremendous heterogeneity, most bidding below their values initially. Bidding rapidly becomes more aggressive, quickly breaking through the threshold of bidding equal to value when this is not imposed as a barrier.[11] Behavior does settle down to a persistent level of overbidding with so low a standard error that even the root prediction squared error is relatively low. The behavior shown in figure 1 is roughly illustrative of the pattern over time in experiments #3–#5, which had somewhat more heterogeneity, and #3 and #4 somewhat higher overbidding, but nonetheless a similar time path. Figure 2 is provided because it contrasts: it shows the only occurrence observed of a substantial time to reach persistent positive mean overbids.

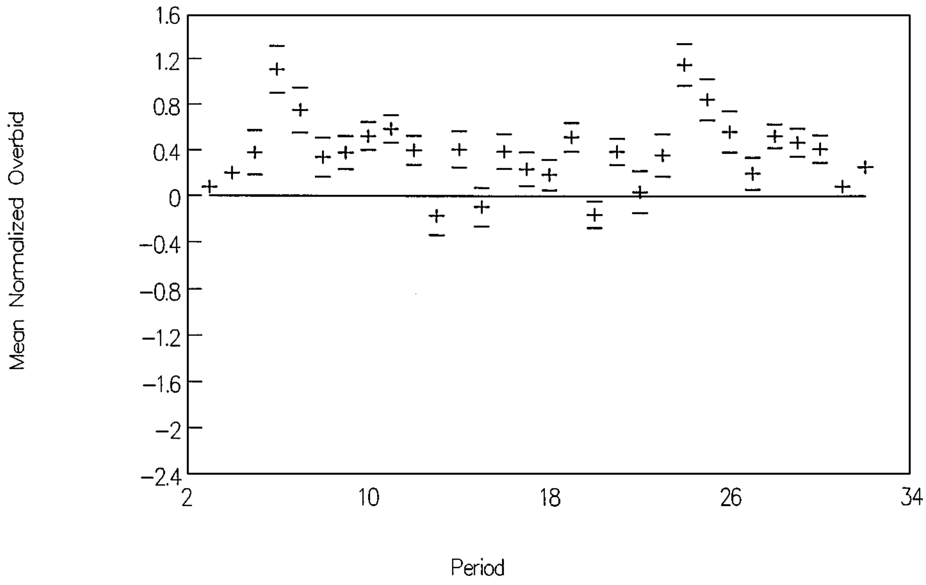

Figure 1.

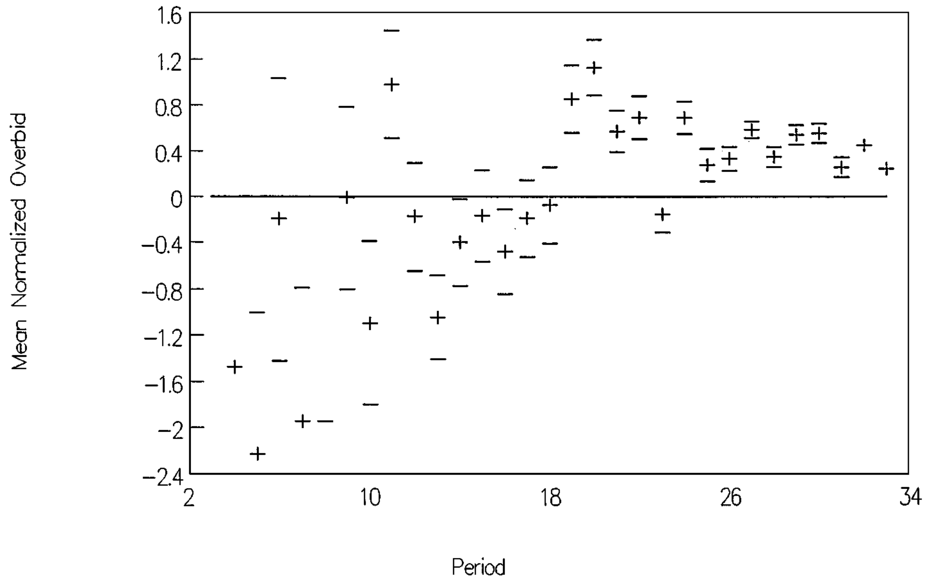

Figure 2.

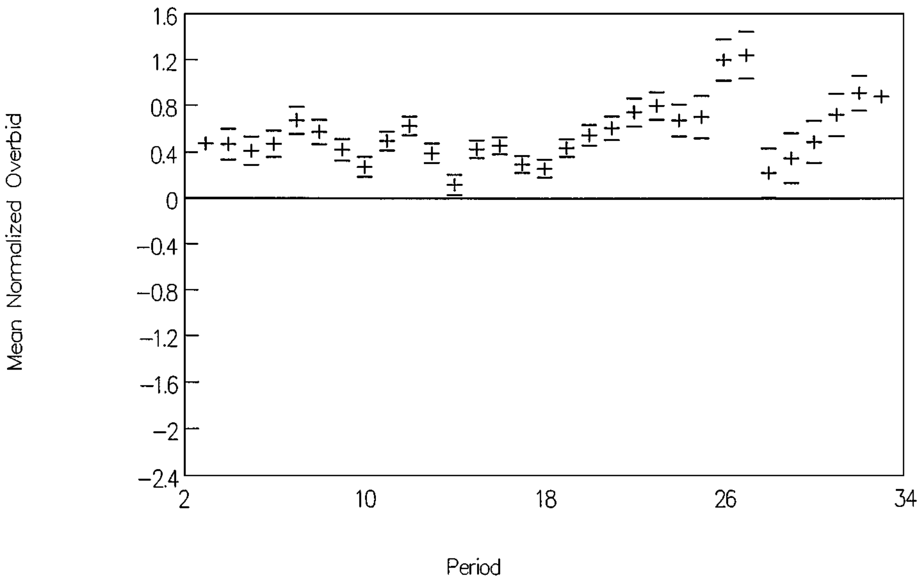

Figure 3.

The underlying individual patterns are largely like the aggregate. Including the experienced subjects in #6, only 3 of 34 subjects (1 each in experiments #2, #4 and #5, with the subject from #2 also taking part in #6) bid their dominant strategy in more than two consecutive periods. All three of these subjects exhibited a stream of dominant strategy choices after 6–11 periods. The feedback provided to other subjects in experiments #2, #4, #5 and #6 indicated that some subject was bidding precisely his value.[12] There is absolutely no sign the competing bidders were at all influenced by this. During informal debriefing, all three subjects who appeared to have adopted the dominant strategy referred to risk aversion as their motivation; none offered any comment even vaguely suggestive of understanding the dominance argument.

B. Comparison to English auctions. Dramatically less deviation from predictions is observed in English auctions, and then almost exclusively in the initial few periods. Prices and mean bidding behavior in English private-values auctions vary less from theoretical predictions than practically any other behavior reported in the laboratory economics literature.

The evidence is also quite strong that the differences between second-price and English auctions will not go away with experience. Learning in English auctions seems to have been completed within the first 4–5 periods. When subjects with 32–33 periods of second-price experience were called back for experiment #6, one subject went bankrupt after 46 periods, and four of the remaining five persisted in overbids averaging almost 0.6, with virtually no remaining variance, through over 75 periods of experience and feedback.[13]

C. The role of feedback. My hypothesis is that the two theoretically isomorphic institutions create differences in feedback. In particular, subjects in second-price auctions may perceive an opportunity to bid more aggressively without necessarily having to pay fully as much as they bid. They are hypothesized to believe that a second-price auction presents a tradeoff, with overbidding leading both to an increased chance of winning the auction and to an increased probability of a loss. They perceive their behavior, in this view, as an appropriate mix of these two attributes, settled upon with experience.

The role of feedback is key to this explanation: the negative feedback when a subject incurs a loss (when both their bid and the second-highest bid exceed their value), if not too frequent or large, is hypothesized to be offset by rounds when a subject overbids and still makes money (because the highest rival bid is below the winner's value), which is viewed (mistakenly) as positive feedback. Note that no occurrence in an English auction could possibly be interpreted in this way: continuing to compete after the price has exceeded a subject's value can only result in losses.

The negative feedback mechanism is quite weak in second-price auctions. To illustrate, consider symmetric bidding at the mean level of overbidding observed in those auctions among #1–#6 with at least five bidders. While expected profit is 57% of predicted profit, the probability that the winning bidder suffers a loss is 0.3147. So under symmetric bidding, the probability that a given bidder loses money in a given auction is 0.0567 (across these data, there are on average 5.55 bidders). Feedback that is hypothesized to be interpreted positively occurs far more often.

This explanation of the observed differences points to feedback in English auctions reinforcing the dominant strategy *every time* the price rises above a subject's value ("If I keep bidding, the only possibility is that I can lose money"), rather than only when a loss occurs. Such a phenomenon appears to be consistent with the psychological literature on feedback and decisionmaking (cf. Camerer, 1995); it certainly serves to organize the observations.

D. Carryover effects of bidders' experience. By itself, nearly 70 periods of experience in second-price auctions appears to be of no help in learning the dominant strategy. Subjects settle into a pattern of substantial overbidding which they may perceive to be an appropriate tradeoff of the frequency of winning versus the odds of suffering a loss.

Experiments #10–#16 observe subjects in second-price auctions who have had a modest amount of experience in another type of auction (11–16 periods of experience). Experiments #10–#12 clearly indicate that this small amount of experience in English auctions serves to reduce significantly overbidding in subsequent second-price auctions.

The conclusion seems unavoidable that subjects *transfer to second-price auctions* some lesson deemed relevant which was *learned in English auctions*, but is *not directly learned* from extensive feedback *in second-price auctions*. Remaining deviations from predictions, while greatly reduced by experience in English auctions, are sufficient to make untenable a view that subjects actually learned the game-theoretic logic precisely (that every time it matters whether you overbid, you necessarily lose money by overbidding). Observations in #13–#16 are designed to help isolate what lessons may be learned by subjects, and carried over to second-price auctions.

E. Price acceptance list auctions. These auctions contain a subset of the features of English auctions that are potentially sources of the reduced overbidding in experiments #10–#12. The highest any unsuccessful bidder is willing to go determines the price, and a discrete grid of prices is considered. The feature that subjects are encouraged to consider each price separately, as a price they are or are not willing to pay (rather than simply consider how high to bid) appears, if anything, to be more transparent in P-List than in English auctions. The missing element is the dynamics of an English auction; those dynamics may serve both to make negative feedback more transparent and to control the order in which subjects ask themselves whether each price is acceptable.[14]

Price acceptance list auctions generate by far the highest variance in overbidding. The mean behavior is an underbid, but dwarfed by the standard error.

Roughly, Table 2 indicates that bidders with experience in P-List auctions overbid by 50% more than those with English experience, but substantially less than those with experience only in second-price auctions. Thus, the features that P-List auctions have in common with English auctions explain part, but not all, of the observed transfer of behavior from English auction experience.

F. First-price auctions. First-price auctions have none of the incentive properties of the other auctions studied in this paper. Subjects should bid less than their value; the actual amount bid is, in theory, determined by a complex tradeoff between probability of winning and profitability given winning, where the terms of this tradeoff depend upon the functional form of each rival's bidding strategy. (The tradeoff which subjects are hypothesized to perceive in second-price auctions does actually arise in first-price auctions.)

Only one feature of first-price auctions comes to mind that might be relevant. If a bidder bids more than his value, only two things can happen: he can be outbid, or he can lose money.[15] It is impossible for overbidding in first-price auctions to obtain the sort of feedback that subjects in second-price auctions are hypothesized to interpret as positive feedback. First-price auctions may perhaps exacerbate, relative to English auctions, the emphasis upon this feature.[16]

Experiments #15 and #16 observed subjects in second-price auctions who had 14–16 periods of experience in first-price auctions. The impact of this experience was both dramatic and highly heterogeneous. Subjects bid, on average, slightly below their value. My anecdotal experience was that these subjects seemed to bid more aggressively over time, as if any learning from first-price auctions about not overbidding might have worn off somewhat with experience in the second-price auctions. However, I could find no significant structural break or significant upward time trend, presumably due in large part to the considerable and persistent heterogeneity across subjects in these data.[17]

A subset of subjects, ranging from 4 to 6 out of 14 (depending upon how stringent a definition is used), overbid in most auctions of #15 and #16 as if no significant lesson was carried over from first-price auctions. Overall, it would appear the simple lesson from first-price auctions, that overbidding yields losses, has a powerful effect upon many second-price bidders, an effect that does not result from second-price feedback alone. It appears that this relatively crude lesson, almost unrelated to the dominance argument, is on the one hand taught to some subjects more powerfully by first-price than by English experience, and on the other a large part of the lesson subjects seem to carry over from English auctions.

4. Nonparametric tests of coefficient restrictions

Goodness-of-fit tests for normally-distributed errors fail. This usual fact of life in human-subject experiments renders the usual parametric methods for hypothesis testing inappropriate. In their place, I test likelihood ratios nonparametrically. Specifically, a null hypothesis is formulated as a set of coefficient restrictions on an estimated linear bid function; the alternative hypothesis is that the coefficient restrictions are not sufficiently likely given the data. The coefficient restrictions are then subjected to likelihood-ratio tests by, maximum-likelihood estimation via the method of Lagrangian multipliers (cf, for example, Kennedy, 1992, pp. 61–72). With this methodology, a transformation of the likelihood ratio is asymptotically Chi-squared, without any assumption of normally distributed errors. Indeed, no parametric assumption on error structures is employed; all that is needed is that errors are drawn from some distribution with finite variance. This nonparametric alternative to T-tests on regression coefficients probably should be more widely used in analyzing experimental data; I am surprised not to have seen this relatively straightforward method used before.

A. Tests. I test two categories of null hypotheses. The first is that the theoretical bid function (bid = value) predicts the observations of bids for a particular type of experience (e.g., second-price bidders with English auction experience), within the class of linear bid functions. That is, the alternative hypothesis is that some linear bid function with a nonzero intercept (and perhaps a slope differing from 1) is sufficiently more likely. The bid functions tested were of the form: $b_{it} = a_1 v_{it} + a_2 R_t + \sum_j a_{3_j} R_t D_{ij} + \epsilon$, where b_{it} is subject i's bid, and v_{it} is his value, in period t, when the radius was R_t. The jth dummy variable D_{ij} is 1 if $i = j$, and 0 otherwise.[18] Thus, the a_2 coefficient represents the intercept in these bid functions, and the a_{3_j} coefficients allow for individual fixed effects.

Table 4 reports results of two specifications of the null hypothesis of no overbidding. The first column pools all data within an auction type and experience class, listing test statistics (minus twice the natural log of the likelihood ratio) for the null hypothesis $a_2 = 0$, versus the alternative $a_2 > 0$, with both null and alternative constrained to $a_{3_j} = 0$. As mentioned, this test statistic is asymptotically Chi-squared. The test statistic of 58.1 (which translates a likelihood ratio of 2.4×10^{-13}) for inexperienced subjects would have less than 1 chance in 200 of being observed if the intercept on an average bid function were approximately 0. The column labeled "Fixed Effects" allows for subject-specific intercepts, testing the null hypothesis $a_{3_j} = 0$ for all subjects i, versus the alternative $a_{3_j} > 0$ for some subjects(s), with both null and alternative constrained to $a_2 = 0$.

The second category of null hypotheses tested are hypotheses that two different types of prior experience have effects on later bidding in second-price auctions that are insignificantly different. The bid functions tested were of the form: $b_{it} = a_1 v_{it} + a_2 R_t + a_3 R_t D_i + \epsilon$, where D_i was 1 for one of the types of experience and 0 for the other (e.g., 1 for subjects who had been in English auctions, 0 for subjects who had only been in second-price auctions). The null hypothesis is the coefficient restriction $a_3 = 0$, versus the alternative $a_3 \neq 0$. Results are reported in Table 5.

Both tables stem from estimated bid functions for which observations in later periods may fail to be independent of observations in earlier periods. This does not seem to be

Table 4. Likelihood of theoretical bid functions, by type of experience (Nonparametric tests of coefficient restrictions).

Pooled data	Significance[a]	Fixed effects	Significance[a]	No. obs.
Second-price auctions, inexperienced subjects				
58.10	(<0.005)	30.88	(<0.005)	520
English auctions				
0.28	Insignificant	21.04	(<0.005)	228
Price acceptance list auctions				
7.46	(<0.025)	9.73	Insignificant	162
Second-price auctions, with experience in:				
—Second-price auction				
114.42	(<0.005)	133.34	(<0.005)	177
—English auctions				
13.75	(<0.025)	47.41	(<0.005)	252
—Price acceptance list auctions				
83.85	(<0.005)	41.04	(<0.005)	120
—First-price auctions				
0.69	Insignificant	19.42	(<0.010)	222

[a]One-sided.

Table 5. Likelihood of comparable effects of experience on later bidding in second-price auctions (Nonparametric tests of coefficient restrictions).

Types of prior experience	Likelihood ratio	Significance[a]
English auctions vs. Second-price auctions	16.48	(<0.005)
Price list auctions vs. Second-price auctions	7.34	(<0.010)
First-price auctions vs. Second-price auctions	24.26	(<0.005)
English auctions vs. First-price auctions	2.47	Insignificant

[a]Two-sided, but the alternative hypothesis is not one-dimensional.

a problem for these particular tests, in that residual errors from estimated bid functions exhibited no significant serial correlations.

B. Discussion. The amount of overbidding in second-price auctions cannot be regarded as potentially due to chance; Table 4 also shows that a hypothesis of no overbidding is wildly implausible for bidders with experience in second-price auctions. In sharp contrast, the likelihood ratio for the pooled observations in English auctions indicates that constraining to a zero intercept and unit slope fits the data almost perfectly.

The coefficient-restriction tests are strongly supportive of a conclusion that extensive experience in second-price auctions does not produce conformity with the theory, and that experience in any of the other types of auctions significantly alters behavior relative to experience solely in second-price auctions. The extreme heterogeneity of the impact of first-price auction experience on behavior in second-price auctions also shows up in Table 4, as the pooled data do not reject the theoretical coefficient restrictions (the test statistic is only 0.69), but data with fixed effects for separate subjects are sharply inconsistent (a test statistic as high as 19.42 has less than 1 chance in 100 if the null hypothesis of no overbidding holds).

The null hypothesis for the fixed-effects tests is extremely narrow, however, in that the alternative includes cases in which a single subject is systematically overbidding.[19] Even the quite clean data from the English auctions is not up to the demands of this null. The failure to reject this null hypothesis for P-List auctions appears to be due largely to an inability of a single intercept to organize individual subject behavior. (A similar phenomenon relates to a relatively low, but still significant, fixed-effects test statistic for subjects with experience in first-price auctions.)

The hypotheses of different types of experience having comparable effects on later bidding in second-price auctions are readily rejected in Table 5. The lone exception is that tests with pooled data cannot distinguish between impacts of prior experience in English and in first-price auctions. The notion that the lesson carried over from P-List auctions is only part of that learned from experience in English auctions is also supported.

5. The challenge to boundedly rational models

These experiments exhibit systematic deviations from rational behavior in second-price auctions. These deviations constitute a challenge to models of boundedly rational behavior, basically for three reasons. First, the rational prediction defied by this behavior is a dominant strategy. Thus subjects' behavior cannot be explained as stemming from boundedly rational beliefs about rivals' behavior or motivations, nor from partial adaptation toward best responses to current or recent actions of rivals. Second, while variability in behavior diminishes with feedback, instead of convergence to rational behavior, I observe its opposite: persistent, indeed increased, overbidding with greater experience and longer feedback. So adaptive models which predict the eventual vanishing of irrational behavior are presumably inconsistent with my observations (though I have not tested what feedback of infinite length might do). Third, while extensive feedback in second-price auctions does not reduce irrationality, a small amount of experience outside the market (in other auctions chosen for the feedback they present) does significantly reduce it.

A variety of bounded rationality models share the characteristic that strategies which perform more favorably come to be used more often. Such models miss the point here: I firmly believe subjects are not comparing the payoffs from the bid functions they employ to the payoffs from bidding their value. The logic that bidding one's value will yield a higher payoff whenever it yields a different payoff does not occur to subjects, and they are not brought to think in this way by any amount of feedback within repeated second-price auctions.

Dominance arguments require visualization. The way in which game theorists articulate their prediction for these auctions ("Suppose it matters whether you bid your value or bid \$5 above your value, . . .") is apparently neither eductively or evolutively natural for subjects. In this light, Bacharach (1993) distinguishes the way in which players label and thereby comprehend the actions they can take.[20] The theorist's notion of a strategy as a mapping from values into bids may not be obvious to laboratory subjects. Even if it is, (i) "Bid truthfully," (ii) "Take no chances," and (iii) "Adopt the dominant strategy" are surely three epistemically distinct descriptions of what happens to be the same mapping. Even when observing a rival whose bids equal her values, a subject still may not perceive any reason to describe her behavior via label (iii) (and, indeed, that may not be apt). He might simply conclude that she is being risk-averse, and decide that his risk tolerance exceeds hers.

To the best of my knowledge, the results reported here defy explanation by any extant bounded rationality model. A survey of that literature is beyond the scope of this paper; the rest of this section illustrates the problems presented to different types of boundedly rational models.[21]

Models emphasizing the psychology of choice tend to focus on the framing of questions as an essential influence on the responses generated (Tversky and Kahneman 1981, 1986, also see Camerer, 1995). While these models have considerable descriptive appeal, they differ from many bounded rationality models in that framing models offer little in the way of constructive structure. Though the failure of these models to come to grips with my observations might be called a technical failure, I think it is fundamental. The question facing subjects in experiments #6–#16 is identical, and identically framed: "Here is your value this period. How much do you bid if the highest bidder wins at the second-highest price?" It is not the context in which this market, the second-price auction, frames its request for a bid, or the way it frames feedback that matters. Rather, it is the way questions and feedback have been framed elsewhere that serves to reduce overbidding. The difference, by the time they get to second-price auctions, is within the subject, not the frame.

One class of models has subjects treat optimization as an exercise that employs scarce mental faculties, and so decisions are suboptimal relative to a fully rational, or, more properly, "free-unlimited-calculating-power" optimum. Instead, errors relative to the fully rational norm diminish as financial incentives for optimization become stronger, or as less costly calculations are required. Such models stem from the work of Herbert Simon. The implication that deviations from perfectly-rational predictions should be evaluated in payoff space was used in Harstad and Marrese (1982), and was popularized by Harrison (1989); its ability to help explain a number of experimental studies was surveyed in Smith and Walker (1993). This whole class of models fails here. Stronger financial incentives in these second-price auctions correspond to larger values of the R parameter, but the data show that absolute overbids are proportional to R—hence absolute errors increase, and relative errors remain constant, as financial incentives increase. Moreover, Smith and Walker (1993) argued that decision cost is systematically lower with experience in the experimental institution, so that the same financial rewards should see subjects coming closer to theoretical predictions as they gain experience. But bidders experienced in second-price auctions adopt larger and more systematic overbids than inexperienced bidders.

A variant on these models is the quantal response model of McKelvey and Palfrey (1995), in which subjects learn with feedback to make fewer and less costly errors, but also seek to take advantage of systematic errors by rival players. While the notion of a quantal response equilibrium helps to explain the data in some experiments, this is not one of them, for the reasons that began this section. The learning seen in sessions #10–#16 above is of an entirely different sort, from an entirely different source.

Finally, Selten's "learning direction theory" (cf. Selten and Buchta, 2000), attempts a more modest goal, simply predicting that subjects whose behavior is more aggressive than the optimum are more likely to switch to a less aggressive than to a more aggressive alternative strategy, and the inverse. Cason and Friedman, (1997), have found that learning direction theory helps explain otherwise anomalous behavior in single call markets. I have tried several econometric specifications similar to theirs, and find no significant impact of learning direction theory in my data, especially in the data for bidders with experience in second-price auctions. Evidence of shifts to bid functions with different intercepts favors shifts to intercepts that are further from the dominant strategy, thus contrary to this theory (though the number of such identifiable shifts is so small as to suggest insignificance). A more straightforward test, also found useful in other contexts where Selten's methods have been employed, is to look at whether subjects bid less aggressively in the period immediately after an observed negative profit (or after their own negative profit), and the inverse. I have also tried a few tests for such next-period reactions, but for none does such an effect show up with a statistical significance approaching 25%. My interpretation is that subjects in second-price auctions believe they face a tradeoff which calls for occasionally incurring losses; there is no sign that observing a loss points to an inconsistency with the way they have optimized in their response to this perceived tradeoff.

In sum, I do not believe that one model of behavior, boundedly or fully rational, can, or indeed should be able to, explain all persistent motivated behavior observed in all contexts economists seek to understand. Several bounded rationality models have increased the scope of laboratory observations which can be explained and perhaps better understood. But for the universe of boundedly rational models to extend their capabilities to predicting the observations here, they must somehow incorporate the role that experience outside the situation has on the perception of strategies and their consequences within the situation. Subjects who have been taught by other experience that overbidding yields losses are the ones who overbid less, or not at all.

6. Concluding remarks

Economic theory has not based a unique prediction of behavior upon fewer or weaker presumptions than for a second-price auction in a private-values environment.

Adoption of dominant strategies, far from being automatic, depends quite specifically on the details of feedback as perceived by subjects. Feedback in English auctions is sufficiently strong and transparent to exhibit thorough and rapid adoption of the dominant strategy. Even persistent feedback in second-price auctions is, it seems, not a very good stimulus for conveying an understanding of dominant strategy arguments. However, a small amount of experience with English auctions gets subjects to view second-price auctions differently, yielding significantly less overbidding relative to the dominant strategy.

Price acceptance list auction experience suggests that a part of what nearly every subject with English auction experience carries over to second-price auctions is a tendency to view bidding as a binary choice, price by price, between willingness to pay, or the price being unacceptably high. Clearly, this is not all the lesson carried over: these data show that the particular dynamics of the English auction also play a significant role. Moreover, for a large subset of subjects, much of the value of English auction experience also appears to result from first-price auction experience. But the logically unrelated first-price auctions appear to have only one crude lesson to offer: thou shalt not overbid.[22]

I have argued elsewhere (Rothkopf and Harstad, 1994) that context matters: to attempt to model an auction market in isolation can limit the model's descriptive and prescriptive usefulness. This is the problem that arises when trying to use antecedent models of boundedly rational behavior. The models would treat a second-price auction in isolation, yet the prior experience bidders have had outside the second-price auction market influences their behavior far more than the feedback within it.

Acknowledgments

Financial support from the National Science Foundation is gratefully acknowledged. John Kagel, Dan Levin, and Douglas Dyer assisted in designing and conducting the experiments. Walter Mayer, Louis Phlips, and Elisabet Rutström offered useful suggestions, and Yung ho Chiu diligent research assistance. Charles Holt, Michael Rothkopf, Ben Sopranzetti, John Wald and two referees assisted to improve the exposition substantially. Any remaining errors are solely the author's responsibility.

Notes

1. The theoretical nicety that a strategy which overbids in second-price or English auctions is only *weakly* dominated offers no comfort here: rivals' bidding in the experiments is sufficiently aggressive that the expected payoff to overbids observed is always strictly and nonnegligibly below the payoff to the dominant strategy.
2. Adoption of *dominated* strategies is observed for Responders in ultimatum games (cf. Crawford, 1997; Roth, 1995, and references they cite). It is clear that a completely different phenomenon is involved: the distinction is transparent to the Responders, who are aware that rejecting small but positive offers leads to a lower payoff. It appears that subjects have privately observed subjective costs of accepting unfair offers which outweigh the pecuniary payoff difference. There is absolutely no hint of a similar limit to laboratory control over subjects' motivations mattering here.
3. Rutsröm (1998) reports a value-elicitation test which is partly motivated by, and consistent with, the results reported here. Harrison, Harstad and Rutström (1995) point to difficulties with assuming that a private-values paradigm can apply to elicited-values settings.
4. After each period, on all subjects' screens, a table listed the bids from highest down to lowest (without identifying which subject made which bid). For most experiments the value of the subject placing that bid was shown to the right of each bid, e.g., the third row might read "3. $87.00, $85.83" if the third-highest bid was $87, and that subject's value was $85.83. In the exceptions, experiments #1, #2 and #4 (as noted in the far right column of Table 1), only the high bidder's value was publicly reported, although all bids were shown. In all experiments, on all screens, the high bidder's value was emphasized, the second-highest bid was emphasized, and the subtraction performed on screen to show the high bidder's profit (or loss). Ties were broken randomly. Trial periods without cash payments began all experiments, with the same parameters as employed for the first period with profits at stake (but different v's). This first actual period is indicated in the column labeled "(periods)" in the table. Experiment #5 set L = H = R in each period, to create statistically

independent private values; experiments #3 and #6 set [L, H] = [$50, $150]; in all other experiments, [L, H] = [$25, $125]. Data for periods 4–18 of experiment #10 were lost in a computer crash. Experiments #1, #2, #7 and #9 were first reported in Kagel, Harstad and Levin (1987).

5. A similar technique has been used independently in Kahneman, Knetsch and Thaler (1990), who ask subjects to indicate whether each of 18 prices is acceptable. Technically, they are not using a dominant strategy auction, as the acceptable price responses are used to generate supply and demand curves, with trading taking place at the price equating these. (In their experiment 5, a randomly determined price corrects this minor incentive difficulty, without any notable change in their observations.) I remark on three differences. They use the approach on both sides of the market simultaneously. Second, their purpose is not to induce values but to discover them, and, accordingly, they do not use neutral instructions: subjects are told to truthfully reveal their valuations. The observations in this paper suggest their instructional alteration has some impact, though not necessarily at a level relevant to their principal conclusions. Third, my subjects had induced values for an abstract commodity unavailable elsewhere; their subjects were bidding to buy or sell commodities with well-established field prices (cf. Harrison, 1994; Harrison, Harstad and Rutström, 1995, for the methodological importance of this distinction).
6. The small amount of data from 3-bidder auctions is excluded from all data analysis. The four-bidder auctions potentially differ. Culling has occurred: two to three bidders have already gone bankrupt. Also, our ex ante confidence in noncooperative behavior is lessened (see Smith, 1982). Actual observations are as follows:

	Mean	Std. Error	Std. Dev.	RPSE	# Obs.
4 bidders:	0.336	0.050	0.596	0.684	144
3 bidders:	0.236	0.045	0.185	0.300	18

No suggestion of tacit collusion arises in the data. As a fraction of R (rather than predicted profit), overbidding in 4-bidder auctions is virtually unchanged from 5- and 6-bidder auctions. This does not completely overcome my reluctance to pool data that ex ante appeared to have reasons for being unpoolable.
7. All pooling done across R and N was supported by non-parametric and Hausman tests, so subjects behave approximately as if this normalization is transparent.
8. Translating into dollars, inexperienced subjects in second-price auctions overbid by $0.94 on average; subjects experienced in second-price auctions, with on average 60% more profit available to them, overbid by $3.11 on average. One subject, representing 17% of the data in experiment #6, never overbid. The remaining experienced subjects overbid $3.83 on average. Variations in N were unavoidable; variations in R effectively increase monetary incentives and improve the power of the statistical tests.
9. In English and Price Acceptance List auctions, predictions incorporate adjustments to the discrete set of bids (i.e., in the v given in the equation). The winner in English auctions was not predicted to have to bid higher than necessary to win.
10. Specifically, the dashes are placed at a distance from the indicated mean equal to one standard error on a 5-period moving average basis. Thus, for example, the dashes corresponding to period 10 indicate the amount of across-subject variability in bidding during periods 8–12.
11. Second-price, independent-private-values auction experiments reported in Coppinger, Smith, and Titus (1980) and Cox, Roberson and Smith (1982) exhibit convergence to the dominant strategy. Their result clearly seems to be due to an artificial constraint imposed by the experimenters, prohibiting bidding above values. That methodology prevents addressing the questions raised here.
12. In experiments #2 and #4, this feedback came only when the truthfully-revealing subject won. However, in experiment #5 and in #6 (when the truthfully-revealing subject from #2 continued this behavior), the feedback, while anonymous, was present every period. Typically, this behavior also stood out in feedback, as all other subjects were bidding in round numbers.
13. Pezanis-Christou (1996) finds essentially the same result, also over 70–75 periods, in two-bidder asymmetric second-price private-values auctions.

14. That is, subjects in P-List auctions chose their own order in which to consider prices. Anecdotally, a couple were observed to consider prices descendingly, and several to consider prices in some non-monotonic order (some starting at their value).
15. Kagel, Harstad, and Levin (1987) report on first-price, affiliated private-values auctions, and hypothesize (footnote 23) that a constraint preventing subjects from bidding above their values had little or no effect. Subjects in experiments #15 and #16 made 288 financially motivated first-price bids without any upper constraint on bids; 287 were below the value.
16. A subject employing game-theoretic reasoning ought to realize that no lesson about overbidding transfers from first-price to second-price auctions, but this observation does not help to understand these data.
17. It may be worth noting that experience in another type of auction does not per se alter behavior: not even slight differences can be found between the first-price auctions conducted before, and those conducted after, second-price auctions in #15 and #16.
18. This formula allows the alternative bid function to differ in intercept or slope, which is technically a different estimation. The distinction is unimportant here. Both versions (restricting only intercept, vs. also restricting slope) were run, obtaining virtually the same likelihood ratios for all datasets considered. Also, testing solely the null hypothesis that the slope of the bid function is 1, without any restriction on the intercept, fails to reject the null on any of the datasets, at significance levels of 0.25 (which is the significance level called "Insignificant" in Tables 4 and 5.) In all tests reported, intercept terms are estimated as multiples of the radius. Throughout, estimating the intercept term as an affine function of the radius yielded insignificant, indeed tiny, differences in likelihood ratios.
19. An alternative hypothesis that some fraction of subjects, say, half, had fixed-effects coefficients that were significantly different from 0 would be preferable for insight. The technology to conduct such a test does not exist.
20. Bacharach's theory of "variable universe games" introduces an element which is seemingly essential to explaining these experiments, whether a strategy ***occurs*** to a player. However, his model as developed only applies to pure coordination games.
21. I have little to say about the models of Crawford (1995, 1997), Fudenberg and Levine (1997), Gale, Binmore and Samuelson (1995), or Roth and Erev (1995), other than that they have proved admirably useful in explaining other experimental deviations from fully rational predictions, and have provided constructive structure, but show no hope that tweaks to their approaches can offer assistance in explaining this second-price overbidding and the impact of experience in other auctions. The problems in applying these models here are those that began this section.
22. It is not even clear why subjects should decide that this lesson remains relevant once the switch to second-price auctions is made. Such a logical failing is also telling.

References

Bacharach, M. (1993). "Variable Universe Games." In K. Binmore, A. Kirman, and P. Tani (eds.), *Frontiers of Game Theory.* Cambridge, MA: MIT Press, pp. 255–275.

Camerer, C. (1995). "Individual Decision Making." In J.H. Kagel and A.E. Roth (eds.), *Handbook of Experimental Economics*. Princeton, N.J.: Princeton University Press, pp. 587–703.

Cason, T.M. and Friedman, D. (1997). "Price Formation in Single Call Markets." *Econometrica*. 65, 311–345.

Coppinger, V.M., Smith, V.L., and Titus, J.A. (1980). "Incentives and Behavior in English, Dutch and Sealed-Bid Auctions." *Economic Inquiry.* 18, 1–22.

Cox, J.C., Roberson, B., and Smith, V.L. (1982). "Theory and Behavior of Single Object Auctions." In V.L. Smith (ed.), *Research in Experimental Economics*, vol. 2. Greenwich, CT: JAI Press.

Crawford, V.P. (1995). "Adaptive Dynamics in Coordination Games." *Econometrica.* 63, 103–143.

Crawford, V.P. (1995b). "Theory and Experiment in the Analysis of Strategic Interaction." Discussion paper 95–37, University of California, San Diego.

Fudenberg, D. and Levine, D. (1995). "How Irrational Are Subjects in Extensive-Form Games?" *American Economic Review*. 85.

Gale, J., Binmore, K., and Samuelson, L. (1995). "Learning to Be Imperfect: The Ultimatum Game." *Games and Economic Behavior.* 8, 56–90.

Harrison, G.W. (1989). "Theory and Misbehavior in First Price Auctions." *American Economic Review.* 79, 749–762.

Harrison, G.W. (1994). "Expected Utility Theory and the Experimentalists." *Empirical Economics*. 19, 223–253. Reprinted in J.D. Hey (ed.), *Experimental Economics*, Heidelberg: Physica-Verlag, 1994.

Harrison, G.W., Harstad, R.M., and Rutström, E.E. (1995). "Experimental Methods and Elicitation of Values," Working paper #B-338, Sonderforschungsbereich 303, University of Bonn.

Harstad, R.M. and Marrese, M. (1982). "Behavioral Explanations of Efficient Public Good Allocations." *Journal of Public Economics.* 9, 367–383.

Hoffman, E., Menkhaus, D.J., Chakravarti, D., Field, R.A., and Whipple G.D. (1993). "Using Laboratory Experimental Auctions in Marketing Research: A Case Study of New Packaging for Fresh Beef." *Marketing Science.* 12, 318–338.

Kagel, J.H., Harstad, R.M., and Levin, D. (1987). "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica.* 55, 1275–1304.

Kahneman, D., Knetsch, J.L., and Thaler, R.H., (1990). "Experimental Tests of the Endowment Effect and the Coase Theorem." *Journal of Political Economy.* 98, 1325–1348.

Kennedy, P. (1992). *A Guide to Econometrics*. Cambridge, M.A: MIT Press.

Krishna, V. and Rosenthal, R. (1995). "Simultaneous Auctions with Synergies," ISP discussion paper #56, Boston University.

McKelvey, R. and Palfrey, T. (1995). "Quantal Response Equilibria for Normal-Form Games." *Games and Economic Behavior.* 10, 6–38.

Milgrom, P.R. (1987). "Auction Theory." In T. Bewley (ed.), *Advances in Economic Theory: Fifth World Congress.* Cambridge: Cambridge University Press.

Pezanis-Christou, P. (1996). "On the Impact of Low Balling: Experimental Results in Asymmetric Auctions." Working paper, University of New South Wales.

Roth, A.E. (1995). "Bargaining Experiments." In J.H. Kagel and A.E. Roth (eds.), *The Handbook of Experimental Economics.* Princeton: Princeton University Press, pp. 253–348.

Roth, A.E. and Erev, I. (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior.* 8, 164–212.

Rothkopf, M.H. and Harstad, R.M. (1994). "Modeling Competitive Bidding: A Critical Essay." *Management Science*. 40, 364–384.

Rutström, E.E. (főrthcoming) "Home-Grown Values and the Design of Incentive Compatible Auctions." *International Journal of Game Theory*.

Selten, R. and Buchta, J. (2000). "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions." In D. Budescu, I. Erev, and R. Zwick (eds), *Games and Human-Behavior: Essays in Honor of Amnon Rapoport*. Mahwah, NJ: Lawren Ehrbaum Associates.

Shogren, J.F., Shin, S.Y., Hayes, D.J., and Kliebenstein, J.B. (1994), "Resolving Differences in Willingness to Pay and Willingness to Accept." *American Economic Review.* 84, 255–270.

Smith, V.L. (1982). "Microeconomic Systems As an Experimental Science." *American Economic Review.* 72, 923–955.

Smith, V.L. and Walker, J.M. (1993). "Monetary Rewards and Decision Cost in Experimental Economics." *Economic Inquiry.* 31, 245–261.

Tversky, A. and Kahneman, D. (1981). "The Framing of Decisions and the Psychology of Choice." *Science.* 211, 453–458.

Tversky, A. and Kahneman, D. (1986). "Rational Choice and the Framing of Decisions." *Journal of Business*. 59 (2), part 2.

Vickrey, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*. 16, 8–27.