7

Auctions: A Survey of Experimental Research

John H. Kagel

Introduction

Auctions are of considerable practical as well as theoretical importance. In practical terms, the value of goods exchanged each year by auctions is huge. In theoretical terms, auctions play a prominent role in the theory of exchange as they remain one of the simplest and most familiar means of price determination in the absence of intermediate market makers. In addition, auctions serve as valuable illustrations of games of incomplete information as bidders' private information is the main factor affecting strategic behavior (Wilson 1992).

In organizing this survey I have relied heavily on Wilson's (1992) and McAfee and McMillan's (1987a) surveys of auction theory. That is, I have chosen to review series of auction experiments that deal with theoretical issues, along with follow-on experiments designed to sort out between competing explanations of the behavior observed. This serves to circumscribe greatly the literature reviewed.

There are two main strands to the literature: private value auctions, where bidders know the value of the item to themselves with certainty, and common value auctions, where the value of the item is the same to everyone, but different bidders have different information about the underlying value. This review is almost exclusively concerned with one-sided auctions, auctions in which there are many buyers and one seller or many sellers and one buyer. Two-sided auctions with many sellers and many buyers are not nearly as well understood theoretically and for this reason have not received the same kind of attention in terms of theory testing as one-sided auctions. (However, see Holt, chapter 5, and Sunder, chapter 6, for the many uses of two-sided auctions in industrial organization experiments and asset market experiments, respectively.)

Part I reviews private value auction experiments. The experimental procedures employed are characterized in section I.A. Section I.B focuses on the Revenue-Equivalence Theorem: In auctions with independently distributed private values,

first-price and Dutch auctions and second-price and English auctions are strategically equivalent, so that prices should be the same in first-price and Dutch auctions and in second-price and English auctions. The revenue equivalence theorem fails as slightly higher average prices are found in first-price compared to Dutch auctions and in second-price compared to English auctions. Experiments aimed at identifying the behavioral basis for these breakdowns are reviewed. Experiments exploring comparative static implications of auction theory—increased bidding in response to increased competition, increased revenue resulting from uncertainty about the number of rival bidders, and increased revenue in response to public information in auctions with affiliated private values—are reported in sections I.C and I.D. Nash equilibrium bidding theory does remarkably well in predicting the sign of these comparative static manipulations. Sections I.E and I.F explore the effect of different types of information feedback on bidding and adjustments in bidding over time (learning) in experimental auction markets. Section I.G evaluates rival interpretations of bidding above the risk neutral Nash equilibrium in first-price auctions, an outcome that is commonly accounted for in terms of risk aversion. The size of this section, a little over a third of section I, is out of proportion to the substantive issue at stake (are bidders really risk averse or does something else underlie bidding above the point predictions of the theory), but this question has sparked considerable controversy among experimentalists. As such it provides a case study in the kind of dialogue and successive interaction that heightens our understanding of behavior.

Part II reviews common value auction experiments. The overriding issue here concerns the existence and persistence of the "winner's curse," when the high bidder ignores the adverse selection problem inherent in winning the auction and winds up paying "too much." Roth, chapter 1, reviews the early experimental work on this topic, which is taken as the starting point for this review. Section II.A reviews studies of the winner's curse in sealed bid auctions with symmetric information. These studies show the winner's curse to be alive and well, at least for inexperienced and moderately experienced bidders. Further, unlike private value auctions, Nash equilibrium bidding theory fails to predict the directional effect of increased numbers of bidders and public information about the value of the item. The winner's curse in English auctions and auctions with asymmetrically informed bidders is explored in section II.B. Experiments investigating the winner's curse in other market settings—bilateral bargaining games with asymmetric information, "blind bid" auctions, and two-sided auction markets where quality is endogenously determined—are reported in section II.C. Section II.D reports on learning and adjustment processes in settings with a winner's curse.

Section III covers "other topics," those that do not fit squarely under sections I and II. Section III.A looks at collusion in auction experiments. Comparisons between field studies and experimental data are offered in section III.B. Sections III.C and III.D conclude with a brief review of two-sided auctions and other applications of auction market experiments.

A brief concluding section summarizes what has been learned to date. I call attention to open research issues in the course of the review.

I. The Symmetric Independent Private-Values Model

The independent private values (IPV) model corresponds to the case where each bidder knows his valuation of the item with certainty and bidders' valuations are drawn independently from each other. Although bidders do not know their rivals' valuations, they know the distribution from which they are drawn. Experimental research has been largely restricted to the case in which valuations, x, are drawn from a uniform distribution $[\underline{x}, \bar{x}]$.

Vickrey (1961) was the first to solve the independent private-values model using a game theoretic formulation. Assuming risk neutral bidders, in a first-price sealed bid auction (in which the high bidder pays the price she bids) the unique risk neutral Nash equilibrium (RNNE) bid function given the uniform distribution $[\underline{x}, \bar{x}]$ is

$$b(x) = \underline{x} + \frac{(n-1)}{n}(x - \underline{x}) \tag{1}$$

where n is the number of bidders in the auction.

The first-price auction is theoretically isomorphic to the Dutch auction where the auctioneer starts with a high initial price and then lowers the price until a bidder accepts the current price. In theory these two institutions yield the same expected price since the situation facing a bidder is the same in both auctions: each bidder must choose how high to bid without knowing the others' decisions and, if she wins, the price she pays is equal to her bid.

In an English auction the price is increased until only one bidder remains. This can be done by having an auctioneer announce prices or having bidders call the bids themselves. The essential feature of the English auction is that bidders always know the level of the current best bid. Here bidders have a dominant strategy of bidding up to their private valuation, x. To bid less than x sacrifices opportunities for winning the item and making a positive profit, no matter how small, while bidding above x and winning results in certain losses. In a second-price sealed bid auction (sometimes referred to as a Vickrey auction) the high bidder wins the item and pays a price equal to the second-highest bid. The bid function here is

$$b(x) = x. \tag{2}$$

This too is a dominant strategy as (i) bidding below x reduces the chance of winning the item with no increase in profit since the second-highest price is paid, and (ii) bidding above x and winning *as a result* of the higher bid results in losses. Note that the dominant bidding strategy in both these auctions does not depend on the number of bidders, risk attitudes, or the distribution from which private values are drawn.

With risk neutral bidders, the *expected* price paid under all four auctions is the same (Vickrey 1961, Meyerson 1981, Riley and Samuelson 1981). This is referred to as the revenue-equivalence theorem. The mechanism underlying this

result is that bids in the first-price auction equal the expected maximum of the distribution of others' valuations (Wilson 1992). The latter corresponds to the predicted dominant strategy price in second-price and English auctions. With risk aversion, the revenue-equivalence theorem breaks down as first-price and Dutch auctions generate greater expected revenue than English or second-price auctions. This follows from the fact that risk aversion promotes bidding above the RNNE in the first-price and Dutch auctions, while the dominant bidding strategy remains unaffected in English and second-price auctions.

Properties of single unit first-price and English auctions extend to multiple unit auctions in which each bidder demands at most one unit of the item and $n > q$, where q is the number of (identical) units of the commodity offered for sale.[1] Uniform price (competitive) auctions in which the k successful bidders pay the highest rejected bid (the k + 1 highest bid) correspond to second-price/English auctions in the sense that each bidder has a dominant strategy of bidding her private valuation. Discriminatory auctions, in which each successful bidder pays the price bid, correspond to first-price auctions in the sense that (i) with risk neutrality, *expected* revenue is the same as the uniform price auction (Milgrom and Weber 1982, Weber 1983) and (ii) with risk aversion, bids will be above the RNNE prediction, yielding greater expected revenue than in the competitive auction (Harris and Raviv 1981).

In what follows I will also have occasion to discuss third-price sealed-bid auctions, auctions in which the high bidder wins the item but pays the third-highest bid price. The third-price auction is a completely synthetic institution, one that does not (and will likely never) exist outside the laboratory. However, to be faithful to Nash equilibrium bidding theory in third-price auctions requires a number of strategic responses quite different from those required in first-price and second-price auctions, examination of which can provide insight into the behavioral processes at work in private value auctions. With bidders' valuations drawn from a uniform distribution, the symmetric RNNE bid function for a third-price auction is

$$b(x) = \underline{x}\,\frac{n-1}{n-2}\,(x - \underline{x}) \tag{3}$$

Note that in the third-price auction $b(x) > x$ for all $x > \underline{x}$ and (somewhat counter-intuitively) $b(x)$ *decreases* for *increases* in n (Kagel and Levin 1993). Further, with constant absolute risk aversion, bidders continue to bid above their private valuations but *below* the RNNE line (and the more risk averse bidders are, the more they bid below the RNNE), the exact opposite of the bidding pattern, relative to the RNNE reference point, in first-price auctions.

A. Experimental Procedures

An experimental session typically consists of several auction periods in each of which a number of subjects bid for a single unit of a commodity under a given pricing rule. The item is often referred to as a "fictitious commodity" or simply as a "commodity," in an effort to keep the terminology as neutral as possible and *not* to relate the auction to any particular market for fear that this may induce subjects to bid in ways they think appropriate to that market (a loss of experimental control). Usually the number of subjects matches the number of active bidders, so that a given set of bidders compete with each other across the different auction periods. This opens up the possibility of super-game effects and bidding rings, issues that are addressed in section III.A below.

Subjects' valuations are determined randomly prior to each auction period and are private information. Valuations are typically independent and identical (iid) draws from a uniform distribution $[\underline{x}, \bar{x}]$, where $\underline{x}$ and $\bar{x}$ are common knowledge. In each period the high bidder earns a profit equal to the value of the item less the price; other bidders earn zero profit for that auction period.[2] Bids are commonly restricted to be nonnegative and rounded to the nearest penny. Some auction experiments, particularly some of the earlier ones, restricted bids to be at or below private valuations (Cox, Roberson, and Smith 1982); the first-price auction series reported in Kagel, Harstad, and Levin (1987). In sealed bid auctions, after all bids have been collected, the winning bid is announced. Some researchers provide additional information: for example Kagel et al. (1987) reported all bids, listed from highest to lowest, along with the underlying resale values and profits of the high bidder (subject identification numbers are suppressed, however), whereas Cox, Smith, and Walker (1988) reported only the high bid. There have been only limited inquiries into the effects of these different information treatments on outcomes (see section I.E below).

English auctions have been implemented using (i) an open outcry procedure where the bidding stops when no one indicates a willingness to increase the price any further (Coppinger, Smith, and Titus 1980) and (ii) an English clock procedure, where prices start at some minimal value and increase automatically, with bidders indicating the point at which they choose to drop out of the auction (which is irrevocable for that period), and the last bidder wins the item at the next-to-last dropout price (Kagel et al. 1987).

In implementing these auctions subjects are sometimes provided with examples of valuations and bids along with profit calculations to illustrate how the auction works (Cox et al. 1985a). Sometimes reliance is placed exclusively on dry runs, with no money at stake, to familiarize subjects with the auction procedures (Kagel et al. 1987). Sometimes both examples and dry runs are employed (Battalio, Kogut, and Meyer 1990). Although there have been no explicit studies of the effects of these alternative training procedures, examples carry with them the danger of implicitly telling subjects how the experimenter expects/wants them to bid, as they provide models which subjects may try to mimic. Dry runs have their problems as well, as subjects may use the opportunity to send signals to their rivals at no cost to themselves.

B. Tests of the Revenue-Equivalence Theorem

Tests of the revenue-equivalence theorem involve two separate issues. The more basic issue concerns the strategic equivalence of first-price and Dutch auctions and of second-price and English auctions. Given the strategic equivalence of

these different auctions, average prices (revenue) are predicted to be the same, irrespective of bidders' risk attitudes. Assuming that strategic equivalence is satisfied, a second issue concerns revenue equivalence between first-price/Dutch auctions and second-price/English auctions. As noted, revenue equivalence here depends on risk neutrality in the first-price/Dutch auctions. However, risk preferences have typically not been controlled for (see, however, section I.G.4 below).

The experimental data show that subjects do not behave in strategically equivalent ways in first-price and Dutch auctions (Coppinger et al. 1980, Cox et al. 1982) or in English and second-price auctions (Kagel et al. 1987). Further, bids in single unit first-price and Dutch auctions are commonly above the RNNE, consistent with risk aversion (section I.G explores the question of risk aversion in some detail). Prices are also commonly above the equilibrium (dominant) strategy prediction in second-price auctions, although here bidding is independent of risk attitudes. In English auctions bidding typically converges to the dominant strategy prediction after a few auction periods.

1. Tests of the Strategic Equivalence of First-Price and Dutch Auctions

Coppinger et al. (1980) and Cox et al. (1982) report higher prices in first-price compared to Dutch auctions, with these higher prices holding across auctions with different numbers of bidders (it is important to explore the effects of varying numbers of bidders here as the bid function depends on n).[3] These results are summarized in Table 7.1. Price differences averaged \$0.31 per auction period, with Dutch prices approximately 5 percent lower, on average, than in first-price auctions. To minimize the effects of between auction variation in prices, which can be sizable, Cox et al. (1982) emphasized paired-comparison designs, auction series in which the underlying distribution of private valuations is the same while the auction institution varies.

Efficiency in private value auctions can be measured in terms of the percentage of auctions where the high value holder wins the item, auctions with Pareto efficient outcomes.[4] First-price and Dutch auctions varied systematically with respect to efficiency levels, with 88 percent of the first-price auctions being Pareto efficient, compared to 80 percent of the Dutch auctions (Cox et al. 1982).

Bidding was significantly above the RNNE in first-price auctions for all $n > 3$, with mean prices in auctions with three bidders being only slightly above the RNNE. Dutch auction prices, while lower than in first-price auctions, were above the RNNE price for all $n > 3$ as well, with Dutch prices in auctions with three bidders somewhat below the RNNE price. Cox et al. (1982) conclude that

> In first-price auctions for groups of size N = 4, 5, 6 and 9, but not for N = 3, we reject the null hypothesis of risk-neutral Nash equilibrium bidding behavior in favor of our version of the Ledyard risk-averse model (where bidders have constant relative risk aversion) of Nash bidding behavior. (33)

Table 7.1. Price Differences: Dutch versus First-Price Auctions

n	First Price	Dutch Price	Difference
3	2.36	1.98	.38
	2.60	2.57	.03
4	5.42	4.98	.44
	5.86	5.68	.18
5	9.15	8.72	.43
	9.13	8.84	.29
6	13.35	13.25	.10
	13.09	12.89	.20
9	31.02	30.32	.70

Source: Cox, Roberson, and Smith 1982, table 7.
Note: Means for paired-comparison auction sequences. Two entries indicate two paired-comparison sequences for that value of n.

In addition, Cox et al. (1982) conjectured that pricing at the RNNE in first-price auctions with three bidders might reflect a breakdown in the noncooperative behavior which underlies the Nash bidding model. However, subsequent experiments report bidding which is substantially, and significantly, above the RNNE in first-price auctions with three bidders (Cox, Smith, and Walker 1988; Dyer, Kagel, and Levin 1989a). Two explanations for these differences immediately suggest themselves. First, Cox et al. (1982) employed a cross-over design with subjects bidding in sequences of Dutch and first-price auctions so that the lower bidding in Dutch auctions may have carried over to lower prices in first-price auctions (see Harstad 1990, reported in section I.B.2, and Rietz (1993), reported in section I.G.4, for reports of hysteresis effects in IPV auctions). Second, expected profit conditional on winning the auction was substantially higher in the later studies, and there is evidence that with higher expected profit bidders act as if they are more risk averse (section I.G.1).[5]

Understanding the Failure of the Isomorphism between Dutch and First-Price Auctions

Cox et al. (1982) offer two alternative explanations for the lower bidding in Dutch versus first-price auctions. One model is based on the assumption that there is a positive utility of "suspense" associated with playing the "waiting game" in the Dutch auction which is additive with respect to the expected utility of income from the auction. The other is a real time model of the Dutch auction in which bidders update their estimates of their rivals' valuations, mistakenly assuming them to be lower than they initially anticipated as a consequence of no one having taken the item as the clock ticks down (what they refer to as the probability miscalculation model).

These alternative formulations are elaborated on and tested in Cox et al. 1983a. The test procedure consisted of tripling the conversion rate from experimental to U.S. dollars in a series of Dutch auctions. Assuming that the utility of money income function has the constant relative risk averse (CRRA) form, and that the utility of suspense associated with the Dutch auction is *independent* of the amount of money involved, tripling the conversion rate will increase the suspense component of the game resulting in higher prices. Contrary to the suspense model's prediction, Cox et al. (1983a) are unable to reject a null hypothesis of no difference in auction prices as a result of tripling the conversion rate.

These results rule out the suspense model specified and are consistent with the probability miscalculation model which predicts no response to the change in conversion rates. However, as Cox et al. (1983a) point out, these no-change results are consistent with a number of other model specifications as well. For example, the suspense model prediction rests critically on the assumption that the utility of suspense is *independent* of the amount of money at stake in the auction. If, as seems plausible, the utility of suspense is a positive function of the amount of money at stake, then whether and how prices will change in the Dutch auction as the conversion rate increases depend on the rate at which the utility of suspense increases with increases in the conversion rate.[6] Further, subsequent research suggests problems with the adequacy of CRRA as a maintained hypothesis and the efficacy of using changes in conversion rates as a treatment variable for determining the effects of incentives on behavior (see section I.G.1 below). Tests based on alternative model specifications, as well as experimental manipulations for which the probability miscalculation model predicts a *change* in behavior, have yet to be conducted. The latter is important since, in general, it is preferable to employ manipulations requiring changes in behavior with respect to a preferred hypothesis in order to feel really comfortable with the validity of that hypothesis.

2. Tests of the Strategic Equivalence of Second-Price and English Auctions

Kagel et al. (1987) report failures of strategic equivalence in second-price and English auctions with affiliated private values. For now we ignore the affiliation issue since the dominant strategy is unaffected by affiliation (affiliation is defined in section I.D along with an analysis of first-price auctions with affiliated private values, in which case affiliation does affect the Nash equilibrium). The results of Kagel et al. for second-price auctions are shown in Figure 7.1. In these auctions prices averaged \$2.00 (11%) *above* the dominant strategy price. These results have since been replicated in independent private value auctions with both experienced and inexperienced bidders (Harstad 1990, Kagel and Levin 1993), so that bidding above the dominant strategy cannot be attributed to affiliation or to subject inexperience. Bidding above the dominant strategy in second-price auctions is relatively wide spread. For example, Kagel and Levin (1993) report that 30 percent of all bids were essentially at the dominant strategy price (within 5 cents of it), 62 percent of all bids were above the dominant strategy price, and only 8

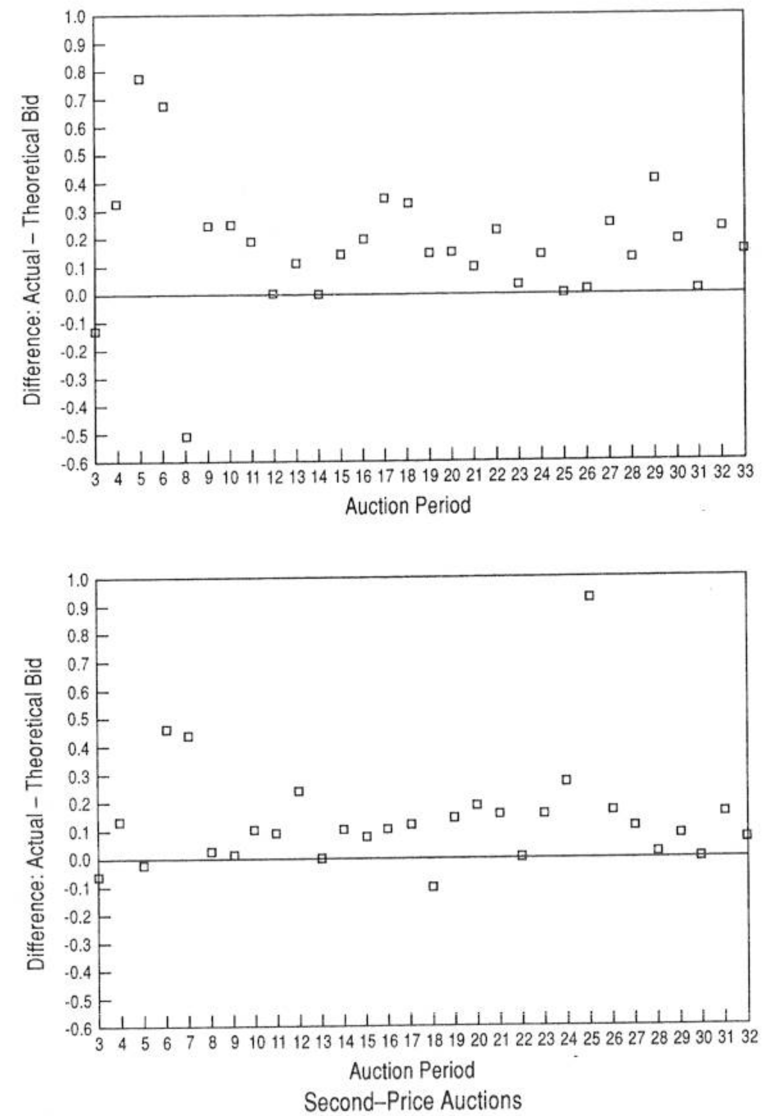

Figure 7.1. Second-price auctions: deviations from dominant strategy price (deviations are normalized by dividing through by the domain from which private valuations are drawn). Top panel, session 1; bottom panel, session 2. *Source:* Kagel, Harstad, and Levin 1987.

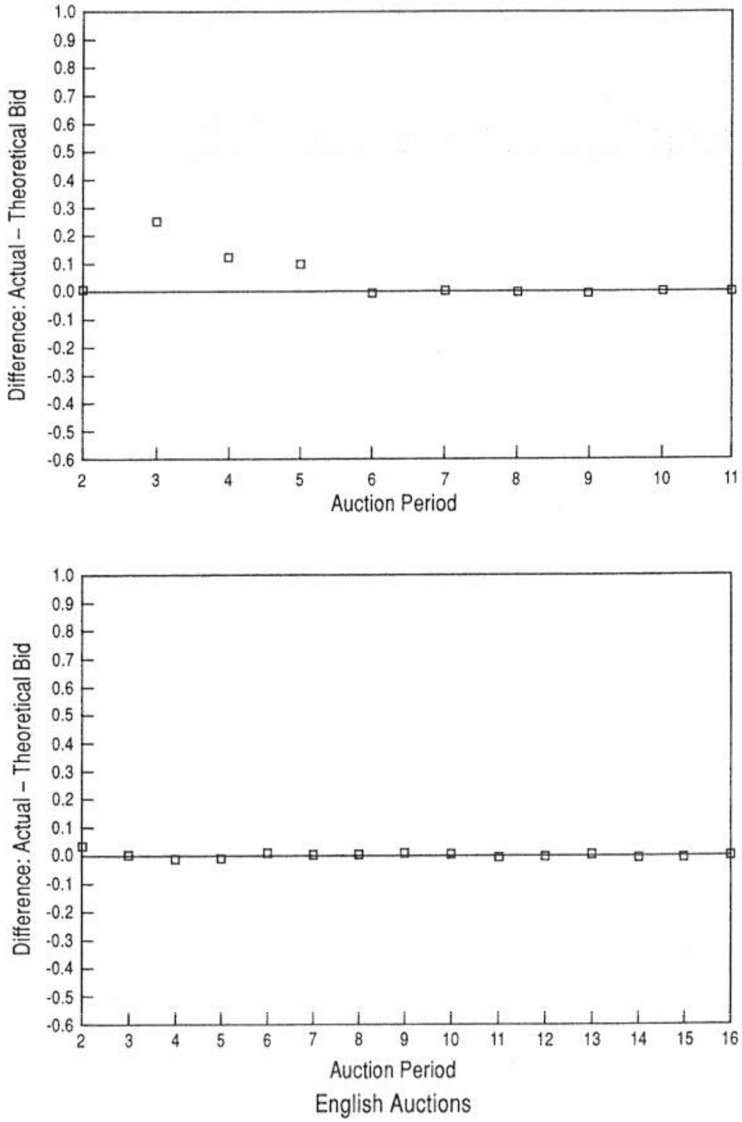

Figure 7.2. English auctions: deviations from dominant strategy price (deviations are normalized by dividing through by the domain from which private valuations are drawn). Top panel, session 1; bottom panel, session 2. *Source:* Kagel, Harstad, and Levin 1987.

percent of all bids were below it.[7] In contrast, as shown in Figure 7.2, in English clock auctions market prices rapidly converged to the dominant strategy price.

Convergence to the dominant strategy price in English auctions is also reported in Coppinger et al. (1980) for single unit IPV auctions and in Van Huyck, Battalio, and Beil (1993) and McCabe, Rassenti, and Smith (1990) for multiple unit IPV English clock auctions (the seller offers multiple units, but each buyer can purchase at most a single unit).[8] Cox et al. (1985a) report that subjects failed to follow their dominant bidding strategy in multiple unit, uniform price, sealed bid auctions. However, unlike the single unit auctions, only a minority of the bids (33 percent for inexperienced subjects, 15 percent for experienced subjects) were above their private valuations, while a majority of the bids were below valuations.[9] Smith and his associates (Smith 1980, Coppinger et al. 1980, Cox et al. 1982) report single unit second-price auctions in which prices converge to the dominant strategy price from below. However, in at least one of these experiments (Cox et al. 1982) subjects were explicitly *prohibited from bidding above their valuations.*

Harstad (1990) reports a variant of the second-price auction (referred to as a "price list" auction) in which subjects indicate whether each of 101 prices (uniformly distributed over the range of conceivable market prices) is acceptable or not, with the bidder with the highest acceptable price winning the item and paying a price equal to the highest of any rival's acceptable price. Although mean prices in these auctions are below the dominant strategy price, these differences are dwarfed by the standard error, indicating a substantial amount of bidding above and below the dominant strategy price. Finally, Harstad (1990) also observed subjects in standard second-priced auctions who had fourteen to sixteen periods of experience in first-price auctions. Mean bids were slightly below the dominant strategy prediction in these second-price auctions, indicating that prior experience in first-price auctions is sufficient to eliminate the average overbidding reported.[10] However, Harstad does note a core of subjects who continuously overbid in the second-price auctions, as if no significant lesson was carried over from the first-price auctions.

The frequency of Pareto efficient outcomes in second-price auctions is quite comparable to first-price auctions. For example, 82 percent of the first-price auctions and 79 percent of the second-price auctions reported in Kagel and Levin (1993) were Pareto efficient.

Understanding the Failure of the Isomorphism between English and Second-Price Auctions

Bidding above the dominant strategy price in second-price auctions would have to be labeled a mistake, since bidding at value is a dominant strategy, irrespective of risk attitudes, the number of rival bidders, the distribution of resale values, and so on. Kagel et al. conjecture that bidding above x is based on the illusion that it improves the probability of winning with little cost as the second-highest bid price is paid. The overbidding observed is sustainable in the sense that average profits are positive. Further, the idea that bidding modestly above x only increases

the chances of winning in cases where you don't want to win is far from obvious. Finally, to the extent that more precise conformity with the dominant bidding strategy results from negative feedback as a consequence of deviating from it, this feedback is weak. For example, for second-price symmetric bid functions of the sort $x + k$, with k equal to the average overbid observed for the high bidders by Kagel et al., the probability of losing money conditional on winning the auction was .36, while the unconditional probability of losing money averaged .06.[11] These punishment probabilities are relatively small, particularly if bidders start with the illusion that modest bidding over valuation increases the probability of winning without materially affecting the price paid; this presumption is usually correct. One straightforward way to test this conjecture is to see what happens to such overbidding when punishment is more likely, as when the number of bidders increase (data reported in section I.C.1 provide somewhat equivocal support for this prediction.)

The structure of the English auction would seem to make it relatively transparent to bidders that they should not bid above their valuations. First, in contrast to second-price auctions, any time you win and bid above your valuation in an English auction, you necessarily lose money. This element of the English auction would appear to play some role because Kagel et al. observed some early overbidding, which collapsed immediately after losses or when losses would have been earned in the dry-run period. Second, the "real time" nature of the English auction would seem ideal for producing observational learning, learning without actually having to lose money, since in comparing the going price with their private value, subjects are likely to see that they will lose money, should they win, whenever the price exceeds their private value. This interpretation is supported by the results of Harstad's (1990) "price list" auctions described in the previous section. Unlike the English auction, in the price list auction you do not necessarily lose money should you win and bid above your private valuation. However, like the English auction, the price list auction focuses attention on whether each price is acceptable or not, which tends to raise the question of whether you want to win with such a bid, much as in the English auction. The net effect is that the difference between bids and values is not significantly different from zero in the price list auction, although there remain a core of bidders who consistently bid above their valuations.

3. Summing Up

The behavioral breakdown of the strategic equivalence of first-price and Dutch auctions and of second-price and English auctions is analogous to the preference reversal phenomenon, where theoretically equivalent ways of eliciting individual preferences do not produce the same preference ordering (see the Introduction, section III.F.1 and Camerer, chapter 8, section III.I).[12] Psychologists attribute the preference reversal phenomenon to a breakdown in procedural invariance principles so that the weight attached to different dimensions of a problem varies sys-

tematically with the response mode employed.[13] In the auctions, prices are higher when bidders must specify a price, as in the first and second-price auctions, compared to the open auctions where the decision is essentially to accept or reject the price that the auctioneer announces. Like the P(rice)-bets in the preference reversal phenomenon, the sealed bid auctions focus attention on the price dimension of the problem, and, like the P-bets, generates somewhat higher prices. On the other hand, the accept/reject decisions involved in the Dutch and English auctions focus attention on profitability, generating somewhat lower prices.

A number of economists have expressed surprise and/or concern regarding subjects' failure to follow consistently the dominant strategy in second-price auctions. The first thing to remember is that the dominant strategy is far from transparent to the uninitiated in second-price auctions (second-price auctions are rarely employed in practice, and Vickrey's seminal paper characterizing the dominant strategy was published in 1961, although economists had, presumably, studied auctions for a number of years prior to this). Note also that there is a small literature by psychologists showing that dominance is often violated in individual choice settings when it is not transparent (Tversky and Kahneman 1986).[14] Further, there is little doubt that if presented with Vickrey's argument, or clear enough examples embodying that argument, subjects would follow the dominant strategy. The latter is, however, quite beside the point in terms of understanding the behavioral forces underlying dominance violations.

In spite of the bidding above valuations observed in second-price auctions, my colleagues and I have argued that the dominant bidding strategy does have some drawing power in explaining behavior. First, one might ask why we don't see even greater deviations from the dominant bidding strategy in second-price auctions? We can only presume that bidders are responding, to some extent, to the forces underlying the dominant strategy: if they bid even higher, the likelihood of winning when they don't want to increases. Second, Kagel and Levin (1993) note that if we declare a bid within $0.05 of a player's private valuation as, for all practical purposes, corresponding to the dominant bidding strategy, then approximately 30 percent of all bids are at the dominant strategy. This is well above the frequency of bidding within $0.05 of private valuations in first-price sealed bid auctions and in third-price auctions, which average less than 3 percent in both cases. Finally, increasing the number of bidders from five to ten, while holding the distribution of resale values constant, does result in higher (more aggressive) bidding in first-price auctions, yet produces no change, or a possible modest reduction, in bids in second-price auctions, which is what the theory predicts (see section I.C.1). (The reduction in bids may result from the increased probability of losing money as a consequence of overbidding with increased n.) Thus, what we are seeing in second-price auctions is relatively stable and modest (in terms of expected cost) bidding above valuations, rather than a complete collapse of the theory. We attribute this overbidding to perceptual errors (response mode effects), given that it does not occur in the English auction, which has the same dominant bidding strategy.

C. *Effects of Varying and Uncertain Numbers of Bidders*

A basic prediction of first-price IPV auctions is that increasing the number of bidders results in higher (more aggressive) bidding. Somewhat counter-intuitively the same treatment effect reduces bids in third-price IPV auctions. Further, if bidders have constant or decreasing absolute risk aversion (which includes CRRA bidders), and there is uncertainty about the number of bidders (numbers uncertainty), this uncertainty raises revenue, on average, compared to revealing information about the number of bidders (McAfee and McMillan 1987b, Matthews 1987). Experiments exploring these comparative static predictions are reported here. In all cases, on average, behavior changes in the way the theory predicts.

1. Effects of Changing Numbers of Bidders

Battalio, Kogut, and Meyer (1990) report a series of first-price IPV auctions in which they exogenously varied the number of bidders while holding the underlying distribution of private valuations constant. Primarily relying on the dual market bidding technique,[15] the large majority of individual subjects (86 percent) increased their bids, on average, when the market size increased from five to ten bidders, with the majority of these increases (60 percent) being statistically significant. Further, none of the subjects who decreased their bids, on average, did so by a statistically significant amount.

Dyer et al. (1989a) report auctions in which bidders made contingent bids, one in a market of size 3 and another in a market of size 6 (in contingent bidding there is uncertainty about the number of bidders in the market so that bids are made conditional on the number of bidders actually present in the market). Seventy-four percent of all contingent bids were higher in the market of size 6 than in the market of size 3. Further, only 3 percent of the bids were in the wrong direction, a larger bid in the market of size 3. For the remaining 23 percent of the observations, the two contingent bids were equal. One reason for the high frequency of equal bids was that subjects did not always change their bids when their valuations were relatively low. These bidders had little chance of winning the auction and therefore little to gain from making the "proper" changes. Discarding the lowest one third of resale values reduced the proportion of equal bids to 12 percent, of the remaining two thirds of the bids, while the error rate associated with bidding more in the small market remained at 3 percent (see section I.G.2 for further discussion of the effects of low expected profit on bidder behavior).

Cox et al. (1988) studied the effects of varying numbers of bidders on the coefficient estimates of individual subject bid functions in IPV first-price auctions. From equation (1) it is clear that the slope of the bid function is increasing in n. The assumption of constant relative risk aversion (CRRA) offered in Cox et al. modifies the bid function as follows

$$b(x) = \underline{x} + \frac{n-1}{n-r}(x - \underline{x}) \tag{4}$$

where r is the coefficient of relative risk aversion (with $r = 0$ corresponding to risk neutrality and $r \leq 1$ required for individual rationality), so that the slope of the bid function is increasing for all $r < 1$. Average sample means of individual subject slope coefficients were increasing for group sizes 3, 4, 5, and 9, with the sample mean for group size 6 below that of group size 4 and 5. Although no pair of means for n = 4, 5, or 6 were significantly different from each other, the mean slope coefficient for n = 3 was significantly less than the rest, and the slope coefficient for n = 9 was significantly more than the rest.

Although it is clear from these first-price auction results that, in general, subjects respond correctly to increased numbers of rivals, it is possible that this is simply a reflex reaction rather than a response to the strategic forces inherent in IPV auctions, since in market settings increased numbers of rivals typically calls for more aggressive behavior. Further, as shown in Kagel et al (1987), there exist relatively simple ad hoc bidding models which predict more aggressive bidding with increased n in first-price auctions. For example, consider the following ad hoc bidding model: Individual i discounts her bid relative to her value by some constant proportion, $\alpha > 0$, of the difference between her value and the lowest possible value, with some accounting for the probable location of rivals in that interval; that is, she uses a discount factor $\alpha(x - \underline{x})/n$. The resulting bid function is

$$b(x) = x - \frac{\alpha(x - \underline{x})}{n} = x + \frac{n - \alpha}{n}(x - \underline{x}) \tag{5}$$

which has all of the qualitative characteristics of the CRRA bid function (4).

The third-price IPV auction provides a more stringent test of the equilibrium prediction of the effects of increased numbers of bidders since it predicts lower (less aggressive) bidding in response to more rivals, contrary to the reflex reaction hypothesis and the ad hoc bidding model's predictions. Kagel and Levin (1993) test this prediction for auctions with five versus ten bidders, with the results reported in Table 7.2. In first-price auctions all bidders increased their average bids, with bids *increasing* an average of \$0.65 per auction period (significant at the 1 percent level). In contrast, in second-price auctions the majority of bidders did not change their bids, with the average change in bids being –\$0.04 (not significantly different from zero). However, in third-price auctions 46 percent of all subjects *decreased* their bids, on average, for an average decrease of \$0.40 per auction period (significant at the 5 percent level). Further, even stronger qualitative results are reported when the calculations are restricted to valuations lying in the top half of the distribution: the average increase in bids in first-price auctions is \$1.22, and 59 percent of all subjects decrease their bids, on average, in third-price auctions, with an average decrease of \$0.86 per auction period. Although a number of bidders in third-price auctions clearly err in response to increased numbers of rivals by increasing, or not changing, their bids, the change in pricing rules has relatively large and statistically significant effects on bidder's responses in the *direction* that Nash equilibrium bidding theory predicts.

Table 7.2. Effects of Changes in Number of Rivals on Bids

	All Private Valuations				Top Half of Private Valuations			
	Changes in Average Bids by Individual Subjects (auctions with $n = 5$ vs. $n = 10$)			Average Dollar Change in Bids (standard error mean)	Changes in Average Bids by Individual Subjects (auctions with $n = 5$ vs. $n = 10$)			Average Dollar Change in Bids (standard error)
Auction	Increased	Decreased	No Change[a]		Increased	Decreased	No Change[a]	
First price	10	0	0	0.65[b] (0.11)	10	0	0	1.22[b] (0.12)
Second price	2	3	5	–0.04 (0.08)	3	1	6	0.02 (0.04)
Third price	16	18	5	–0.40[c] (0.17)	11	23	5	–0.86[b] (0.26)

Source: Kagel and Levin 1993.
[a]Average absolute difference is less than or equal to $0.05.
[b]Significantly different from 0 at 1% level.
[c]Significantly different from 0 at 5% level.

2. Uncertainty Regarding the Number of Bidders

In most theoretical auction models the number of competing bidders is assumed to be fixed and known to all bidders. McAfee and McMillan (1987b) and Matthews (1987) generalize auction theory to allow uncertainty regarding the number of bidders (numbers uncertainty). They show that in IPV first-price auctions if the number of bidders is unknown and bidders have constant or decreasing absolute risk aversion—a belief that most economists hold as a working hypothesis (Arrow 1970, Machina 1983) and which includes CRRA—then expected revenue is greater if the actual number of bidders is concealed rather than revealed.[16] In contrast, if bidders are risk neutral, expected revenue is the same whether the actual number of bidders is revealed or concealed. Examining the effects of numbers uncertainty provides an internal consistency test of the hypothesis that bidding above the RNNE in first-price auctions reflects, at least in part, elements of risk aversion.

Dyer et al. (1989a) compared a contingent versus a noncontingent bidding procedure. The contingent bidding procedure served to reveal the number of bidders by permitting each bidder to submit a vector of bids, with each bid in the vector corresponding to a specific number of rivals, and with that bid binding when the realized number of rivals matched the number of the contingent bid. In

the noncontingent bidding procedure only a single bid was permitted and had to be made prior to the random draw determining the number of active bidders in the market, thereby concealing information about the number of rivals. The market revenue predictions of the theory were clearly satisfied as the noncontingent bidding procedure raised more revenue than the contingent procedure, with a mean difference of $0.31 (which is significant at the 1 percent level) on an average revenue base of $20.67. Although the revenue predictions of the theory were satisfied, examination of individual bids showed that (i) narrowly interpreted, the Nash equilibrium bidding theory underlying the market predictions was rejected, as less than 50 percent of all bids satisfied the strict inequality requirements of the theory, but (ii) a majority of the deviations from these inequality requirements favored the revenue-raising predictions of the Nash model, and, in a large number of cases, involved marginal violations of the theory.

Dyer et al. (1989a) score these results as a partial success for the theory and the argument that bidding above the RNNE in private value auctions is, at least in part, a result of risk aversion. This evaluation is based on the evidence reported and on the fact that the theory assumes symmetry (identical bid functions) and that all bidders, no matter how small the expected gain, will find it worthwhile to adjust their bids, assumptions which are unlikely to be strictly satisfied, so that it is probably asking too much to expect the point predictions of the theory to be exactly satisfied. I return to the question of risk aversion in private value auctions in section I.G.

3. Summing Up

Examining the effects of varying numbers of bidders tests one of the basic comparative static implications of auction theory. Increased numbers of rivals almost always results in higher (more aggressive) bidding in first-price auctions, at least for bidders who perceive themselves as having a realistic chance of winning the auction. This is not just a reflex reaction to increased competition, as the same treatment effect has essentially no impact on bidding in second-price auctions and results in lower bids in third-price auctions, as Nash equilibrium bidding theory predicts. With numbers uncertainty, concealing (as opposed to revealing) information about the number of rivals raises average market prices, as predicted if bidders exhibit constant or decreasing absolute risk aversion. This last result serves as qualified support for the idea that bidding above the RNNE in first-price auctions results, in part, from risk aversion.

D. Auctions with Affiliated Private Values

In auctions with affiliated private values, bidders know the value of the item to themselves with certainty, but a higher value of the item for one bidder makes higher values for other bidders more likely (private values are positively correlated relative to the set of possible valuations).[17] Auctions with affiliated private values have a number of comparative static implications with interesting

behavioral and normative implications. In particular, the effects of public information about rivals' valuations should, on average, increase revenue assuming risk neutral bidders (the impact of public information about item valuation is particularly prominent in common value/mineral rights auctions and will be discussed in detail in section II.A). There are clear predictions regarding individual responses to such information which can be readily tested as well.

Kagel et al. (1987) used a two step procedure to generate affiliated private values: First, a random number x_o was drawn from a uniform distribution on $[\underline{x}, \bar{x}]$. Second, once x_o was determined, private values x were determined, one for each bidder, randomly drawn from a uniform distribution $[x_o - \epsilon, x_o + \epsilon]$, where ϵ was common knowledge. Announcing x_o provides maximum public information about the distribution of rivals' private valuations, reducing the auction to one with IPVs.

For private valuations in the interval $\underline{x} + \epsilon \leq x \leq \bar{x} - \epsilon$, the RNNE bid function when x_o is not known is

$$b(x) = x - \frac{2\epsilon}{n} + \frac{Y}{n} \tag{6}$$

where Y is a negative exponential which becomes negligible rapidly as x moves beyond $\underline{x} + \epsilon$. With x_o announced, this reduces to the following IPV bid function:

$$b(x, x_o) = (x_o - \epsilon) + \frac{n-1}{n}[x - (x_o - \epsilon)] \tag{7}$$

(the lower bound of the interval, $x_o - \epsilon$, serves the role of $\underline{x}$ in the standard IPV auction design—recall equation (1)). Under the RNNE, revealing x_o raises individual bids unless x is very close to $x_o + \epsilon$. The intuition here is straightforward: In symmetric bidding models each agent bids as if she has the high signal value since this is when she wins, the only event that counts. As such, bidders with relatively low valuations are, on average, surprised following the announcement of x_o, realizing that they had lower valuations than they had assumed prior to the announcement, and raise their bids in efforts to win the item. This in turn puts pressure on bidders with higher valuations to raise their bids, resulting in increased prices and more revenue.

With affiliated private values and risk aversion, bids are above the RNNE model prediction. Further, with risk aversion, the release of public information, even large amounts of public information, such as announcing x_o, may not result in increased revenue (Milgrom and Weber 1982, Kagel et al. 1987).[18] However, assuming symmetry in bidding, Kagel, Harstad, and Levin demonstrate that for any concave utility function, all bidders whose private valuations lie below $C(x_o) = x_o + [(n-2)/n]\epsilon$ will raise their bids following the announcement of x_o. This prediction has considerable bite since on average the highest valuation lies just above $C(x_o)$.

Kagel et al. (1987) contrast the Nash equilibrium bidding model's predictions with two ad hoc bidding models. In the ad hoc models bidders do not consider the best response to their rivals' behavior, but simply discount their bids taking account of rudimentary strategic forces inherent in varying ϵ and n. One might justify bidding schemes of this sort with reference to the strong informational and cognitive requirements inherent in calculating a best response to rivals' behavior. As such, the ad hoc models serve as benchmarks against which to evaluate bidders' responsiveness to the strategic forces underlying the more sophisticated Nash equilibrium bidding theory.

The ad hoc bidding models can readily account for bidding above the RNNE in auctions with affiliated or independent private values, and predict linear bid functions (section I.C.1, equation (5), formulates one such model for the IPV case). But the ad hoc models also predict that revealing x_o will result in no change, or a decrease, in average winning bids. Further, in the ad hoc models for all $x > C(x_o)$ bids will be lowered following the announcement of x_o, in contrast to the Nash model's prediction that bids are likely to increase in this interval.

Bidding was studied under several values of ϵ. With $\epsilon = \$6$, the smallest ϵ value employed, high bids (prices) averaged \$0.29 *below* the RNNE prediction (these deviations are significant at the 10 percent level).[19] In contrast, with $\epsilon = \$12$ and \$24, prices were significantly *above* the RNNE prediction (averaging \$1.40 and \$3.34 above it). Both the ad hoc and the Nash bidding model with risk aversion can account for these data, but to do so the markdown coefficient in the ad hoc model would have to vary with ϵ and the utility function underlying the Nash model would have to exhibit increasing, as opposed to constant, relative risk aversion.[20]

Public information about rivals' valuations (announcing x_o) increased prices an average of \$0.22 per auction, which is about 30 percent of the increase predicted under the RNNE and is not significantly different from zero. Further, these changes in revenue were quite variable across auctions (substantially more so than equilibrium predictions), with most of the increase being accounted for by one of five auction series. However, public information announcements were only implemented with $\epsilon = \$12$ and \$24, when bidding under private information conditions was well above the RNNE. As such, the results are closer to the predictions of a risk averse Nash equilibrium bidding model that permits some increase in revenue, than the ad hoc models which predict no change or a decrease in prices with the release of public information.

Table 7.3 shows how individual subjects altered their bids following the announcement of x_o. In cases where $x < x_o$ (the first column in Table 7.3), both the ad hoc and the Nash equilibrium bidding models predict that subjects will raise their bids. This happened 67 percent of the time, with the large majority of the deviations from this prediction involving no change in bids. Undoubtedly, a partial explanation for the relatively large numbers of unchanged bids here is, as one subject remarked, a result of the slim chances of winning the auction, so that it did not pay to bother calculating a bid increment (similar results were reported for bidders with low resale values in auctions with varying numbers of bidders, section I.C.1).

For bidders with resale values in the interval $x_o < x < C(x_o)$, both the sophisticated ad hoc bidding model and the Nash equilibrium bidding model predict that

Table 7.3. Effects of Public Information on Individual Bids with Affiliated Private Values

Bids	Private Values Relative to x_0		
	$x < x_0$	$x_0 \leq x \leq C(x_0)$	$x > C(x_0)$
$b(x,x_0) > b(x)$	66.8%	67.0%	38.2%
$b(x,x_0) = b(x)$	27.3%	17.0%	18.2%
$b(x,x_0) < b(x)$	5.9%	16.0%	43.6%
Total number of bids	187	106	55

Source: Kagel, Harstad, and Levin 1987.

Notes: Each entry is the percentage of the values drawn, relative to x_0, as in the column heading, for which the subject responded as in the row heading. Analyses restricted to $\underline{x} + \varepsilon \leq x \leq \bar{x} - \varepsilon$. $C(x_0) = x_0 + (n-2)\varepsilon/n$.

bids will increase with the release of public information. This happened in 67 percent of the cases, well above the frequency expected by chance factors alone. Attempts to rationalize the failure to change bids here on grounds of subjective transactions costs, in conjunction with a low probability of winning the auction, are on substantially weaker footing than when $x < x_o$, and should be counted as violations of both the ad hoc and the Nash equilibrium bidding models.

The last column in Table 7.3 shows what happens to bids when $x > C(x_o)$. Nash equilibrium bidding theory, accounting for risk aversion, calls for increasing or decreasing bids here, depending on the degree of risk aversion and the location of private values in the interval $[C(x_o), x_o + \epsilon]$. In contrast, even the most sophisticated of the ad hoc bidding models calls for *all* bidders reducing their bids here. Instead, a sizable proportion increase their bids, consistent with the Nash bidding model's prediction.

Summing Up

In first-price auctions with affiliated private values, Nash equilibrium bidding theory organizes the data better than either of two ad hoc bidding models. Public information about others' valuations increases average market prices, but the increase is smaller and less reliable than predicted under RNNE bidding. Lower average revenue might be attributed to risk aversion, and the high variability may be attributed to the sizable frequency of individual subject bidding errors relative to the theory.

E. Effects of Price Information in Private Value Auctions

There is considerable diversity in price information feedback following submission of sealed bids in field environments and between different research groups in the laboratory. For example, government agencies typically report back the full set of bids, along with bidders' names, as they are required to do so by law, but the private sector typically does not report any prices to losing bidders, just that they were not successful. Although a persistent bidder can often obtain fairly complete price information on a private sector job, most do not, and for those who do the data are less precise than on public sector bids.[21]

Isaac and Walker (1985) have studied the effect of price information feedback in IPV auctions. Under their full information condition subjects were provided with all bids from the previous period and the bidder's identification number. In the limited information condition subjects were given only the winning bid from the previous period and the identification number of the winner. In first-price auctions with four bidders, prices under the limited information condition were consistently higher than under the full information condition, but there were no significant differences in efficiency between the two information conditions.

Standard noncooperative game theoretic models of auctions make no mention of the effects of price information following bid submission since most theory relates to single period auctions where such feedback would be irrelevant.[22] It is interesting to speculate as to how price feedback impacts on bidding. In the auction experiments, unlike the single-period theory, a fixed number of players bid in a series of twenty-five auction periods. This suggests two possibilities (which are not mutually exclusive): First, the full information condition may provide bidders with a chance to signal their intentions not to bid very aggressively when they have low valuations. That is, the added information promotes tacit collusion to some extent (although winning bids are still above the RNNE prediction). Second, bidders are unlikely to determine immediately what their best strategy is and to have to go through a trial and error learning period as they explore the effects of different bidding strategies. In this case the price information may provide them with a better sense of what their rivals' bids are likely to be and to help correct for overly pessimistic, or overly optimistic, initial expectations. (Section I.F briefly examines adjustment/learning patterns in private value auctions.)

Cox, Smith, and Walker (1984) examined the effects of suppressing information about the highest accepted and highest rejected price (their standard price information feedback) in multiple unit discriminatory price auctions. They found no effect on bidding. Battalio et al. (1990) examined the effect of reporting the winning bid versus the highest three bids in auctions with five and ten bidders. They found no systematic response to the increased information. The suggestion is that relatively minor variations in price information, compared to the major variation implemented in Isaac and Walker (1985), have little effect on bidding.

Finally, Isaac and Walker also studied the effects of information in auctions where bidders were permitted to discuss strategy prior to bidding and to form cartels. These results will be reported in section III.A, which deals with collusion in auctions.

F. Learning, Adjustment Processes, and Cash Balance Effects in First-Price Private Value Auctions

There has been very little study of learning and adjustment processes in private value auctions as, eyeballing the data, investigators typically observe no pronounced trends in bids relative to resale values.[23] There is, however, increasing

Table 7.4. Time Trend and Cash Balance Effects in Private Value Auctions

Isaac and Walker (1985)[a]	
Full Information:	
$0.826\,x - 0.246\text{ peri} + 0.003\text{ bal}_{t-1}$ $(0.006)^b \quad (0.102)^c \quad (0.008)$	$R^2 = .954$
Limited Information:	
$0.858\,x - 0.530\text{ peri} + 0.012\text{ bal}_{t-1}$ $(0.007)^b \quad (0.122)^b \quad (0.014)$	$R^2 = .938$
Battalio, Kogut, and Meyer (1990)[d]	
$0.936\,x + 0.026\text{ dx} + 0.022\text{ peri} + 0.008\text{ bal}_{t-1}$ $(0.004)^b \quad (0.003)^b \quad (1.04) \quad (0.013)$	$R^2 = .979$
Kagel, Harstad, and Levin (1987)[e]	
$0.984\,x - 0.063\varepsilon - 2.50\text{ peri} - 0.031\text{ bal}_{t-1}$ $(0.002)^b \quad (0.014)^b \quad (0.635)^b \quad (0.023)$	$R^2 = .996$

Notes: peri = 1/(auction period); bal_{t-1} = cash balance at the time of the bid; dx = change in slope coefficient for $n = 10$. Numbers in parentheses are standard errors.

[a] ($n = 4$, $\bar{x} = \$10$, IPV).

[b] Significantly different from 0 at 1% level, two-tailed t-test.

[c] Significantly different from 0 at 5% level, two-tailed t-test.

[d] ($n = 5$ & 10, $\bar{x} = 29$–59, IPV).

[e] ($n = 6$, affiliated private values).

evidence of systematic adjustments in bidding over time in first-price auctions. Smith and Walker (1993), using mean deviations from RNNE bidding as the dependent variable, report that more experienced bidders (those participating in their second, third and fourth auction series) bid significantly higher, as if they are more risk averse, than inexperienced bidders in auctions with four bidders.

Dyer (personal communication) in a limited inquiry, used a fixed effect regression model to identify relatively large, systematic reductions in bids over time for inexperienced bidders in auctions with three bidders.[24] Dyer found further reductions in bids when these subjects participated in a second and third auction series, with the adjustment process essentially completed by the end of the second auction series. These results, in conjunction with Smith and Walker's, suggest that the most pronounced adjustments in bids will occur for inexperienced bidders and that a nonlinear time trend will provide a better fit to the data than a linear adjustment process. Finally, in Dyer's experiment adjustments in bids were sufficiently strong that the coefficient value for the resale variable (x) dropped from 0.764 for inexperienced bidders to just barely above the RNNE prediction of 0.666 for experienced bidders.

Further exploration of adjustments in bidding over time are reported in Table 7.4

where fixed effect regression estimates of time trends in bidding are shown for several different private value auction experiments.[25] In each case the time trend variable is the inverse of the auction period, involving a nonlinear time trend specification. In the Isaac and Walker experiment the time trend coefficient is significant under both full and limited information conditions, with the negative coefficient value indicating *higher* bidding over time. Interestingly, the data suggest a more pronounced increase in bids under limited, as compared to full, information conditions. The time trend coefficient is not significant in the Battalio et al. experiment. Finally, the Kagel et al. experiment, with affiliated private values, shows a pronounced increase in bidding over time as well.

Specifying a nonlinear, as compared to a linear adjustment process seems superior on theoretical grounds as it allows more rapid learning initially and permits the adjustment process to converge to some steady state behavior rather than continuing forever at the same rate. It also tends to generate more significant results, as a linear time trend specification fails to achieve statistical significance at conventional levels for the Kagel et al. data and for the Isaac and Walker data under full information conditions and is only marginally significant (at the 10 percent level) in the Isaac and Walker data under limited information conditions. Taken together with Dyer's results, these data suggest that, more often than not, there is some initial adjustment in bids in private value auction experiments, and that the speed and size of adjustment (and even its direction) may be sensitive to the number of competing bidders.

In an effort to capture the impact, if any, of prior earnings on bids, the lagged value of subjects' cash balances (bal_{t-1}) are also included in the regressions reported in Table 7.4. With constant absolute risk aversion, or with subjects evaluating outcomes in terms of deviations from the current status quo in each auction period (as in prospect theory, Kahneman and Tversky 1979), bids should be independent of cash balances. This is a necessary condition for treating these auction series as a collection of single shot auctions. Consistent with this requirement, the cash balance variable does not achieve statistical significance at anything approaching conventional levels in any of these experiments.

G. Risk Aversion, CRRA, the Flat Maximum Critique, and the Binary Lottery Procedure for Controlling Risk Preferences

Bidding above the RNNE outcome is the most common outcome in single unit first-price private value auctions. (Figure 7.3 provides representative data on this score for high bidders in auctions with $n = 3$ and 6.) Bidding above the RNNE outcome can be rationalized in terms of risk aversion. This insight has generated a number of research results, as well as considerable controversy among experimentalists. This section reviews these results and the controversies surrounding them in some detail, as it provides a case study in the give and take of experimental inquiry that results in increased understanding of behavior.

In response to the observation that bidding is typically above the RNNE in

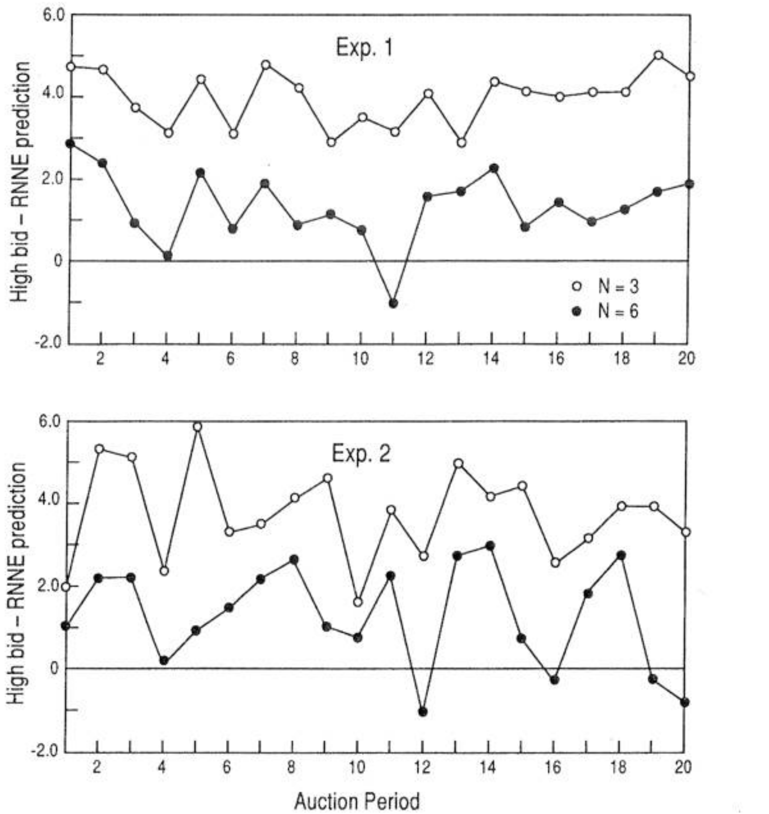

Figure 7.3. First-price auctions: deviations from RNNE price. *Source:* Dyer, Kagel, and Levin 1989a.

single unit first-price auctions, Cox, Smith, and Walker developed a model of heterogeneous bidders with constant relative risk averse (CRRA) utility functions (hereafter referred to as CRRAM; see Cox et al. (1988) for a summary of their model and results). CRRAM provides a focus for much of what follows since Cox, Smith, and Walker have vigorously pursued its development and actively defended it against alternative explanations. The objective of CRRAM is to provide a unified account of bidding above the RNNE in first-price private value auctions in terms of the maintained hypotheses of risk aversion and equilibrium bidding. This has led to a number of criticisms.

First, CRRAM fails as a maintained hypothesis even in terms of characterizing risk aversion in first-price auctions. Second, Harrison (1989) has argued that Cox, Smith, and Walker's conclusions regarding risk aversion are not well supported, as the expected cost of deviating from RNNE bidding is quite small (less than \$0.05 at the median), so that in terms of expected monetary payoffs ("payoff space") many subjects had little to lose from such deviations (the flat maximum

critique). This has set off a vigorous debate among experimentalists (see the September 1992 issue of the *American Economic Review*).[26] Third, although risk aversion organizes bidding in some private value auction environments, it fails in others. This casts doubt on risk aversion as the primary causal factor behind bidding above the RNNE in first-price auctions. Fourth, the failure of the binary lottery technique (Roth and Malouf 1979) to eliminate higher than predicted bids in first-price private value auctions calls into question the joint hypotheses of risk aversion *and* equilibrium bidding (however, results here are sensitive to the way the technique is introduced and the statistical procedures used to evaluate it).

Results of this dialogue suggest that with respect to the primary issue, the role of risk aversion in first-price auctions, it is probably safe to say that risk aversion is one element, but far from the only element, generating bidding above the RNNE. This is not to say that there is no longer any debate regarding the relative importance of risk aversion versus other factors, or in terms of what these other factors are.

1. Risk Aversion and CRRAM As Applied to Single Unit First-Price Auctions

Market prices from first-price private value auctions generally exceed the RNNE prediction irrespective of the number of bidders in the auction or the research group conducting the investigation (sections I.B.1 and I.C). Examining individual subject bidding data from their own auctions, Cox et al. (1988) measure deviations from the RNNE outcome. Applying a Wilcoxon signed ranks test to these deviations, for 91.1 percent of the subjects the deviations lie above the RNNE (in 75 percent of these cases the deviations are significant at the 5 percent level using a two-tailed test). Thus, the direction of deviations from the RNNE are overwhelmingly on the side predicted for risk averse bidders. Employing a second nonparametric test to these deviations, for 60 percent of the auction series (twenty-eight out of forty-seven), a null hypothesis of identical bid functions is rejected. Cox et al. (1988, 73) conclude, "Bid diversity is a prominent, but not extreme, characteristic of our subject pool."

The CRRAM hypothesis implies that each agent's bids will be a homogeneous linear function of resale values for bids that do not exceed an upper bound, defined as the maximum bid that would be entered by the least risk averse bidder in the population (b^*).[27] This leads to the regression specification

(8) $$b_i = \alpha_i + \beta_i x_i \quad \text{for } x_i \le \frac{n - r_i}{n - 1} b^*$$

where we have suppressed the random error term, all bids and valuations are measured relative to $\underline{x}$, and the subscript i refers to the fact that α, β, and r have subject specific values. CRRAM implies that $\alpha_i = 0$ and $\beta_i = (n-1)/(n-r_i)$, where r_i is the coefficient of relative risk aversion for bidder i (see equation (4), section I.C). Fits of (8) to individual subject data yield uniformly high R^2 values of 0.96 or better for 80 percent of the individual subjects (Cox et al. 1988). However, for

22 percent of these bid functions, the intercept term, α, is significantly different from zero. Further, α is not randomly distributed around zero, with 63 percent of $\alpha < 0$, more than would be expected on the basis of chance factors alone. These significant intercept values indicate that "some bidders bid zero or a smaller fraction of value at low values or that other bidders bid consciously and deliberately in excess of value." (Cox et al. 1988, 79–80).[28] Cox et al. (1988) have also tested CRRAM against a general log-concave model by adding squared and cubed values of x_i to (8). For 38 percent of the individual subject data tested, they reject the null hypothesis that the bid functions are strictly linear. These results indicate that CRRAM forms a good, but far from perfect, characterization of Cox et al.'s (1988) individual subject bidding data. See, however, Rietz (1993, section I.G.4) who argues that failure to account for heteroskedastic errors and censoring of bids biases the standard errors of the regression coefficients downward, resulting in larger type I errors than the assumed value.

Tests of CRRAM

A unique implication of CRRAM is that multiplying the profit of a winning bid by any factor $\lambda > 0$ affects only the utility scale, but has no effect on the equilibrium bid (Cox et al. 1983b, 1988). To test this prediction Cox et al. (1983b) compared market prices from first-price auctions with a one-to-one conversion rate of experimental to U.S. dollars to auctions with a three-to-one conversion rate. Cox et al. (1988) later compared mean coefficient estimates of α_i and β_i from individual subject bid functions for these same data. In both cases they fail to reject a null hypothesis of no change, consistent with the CRRAM model.

In contrast, other tests yield less favorable outcomes for CRRAM. Kagel et al. (1987) (section I.D) found that as expected profit conditional on winning the auction increased, bidders increased their bids proportionately more, earning a smaller share of profits compared to the RNNE prediction, a result that is inconsistent with constant relative risk aversion, but which is consistent with increasing relative risk aversion.[29] Since, within the context of CRRAM, these results could possibly be explained by the nonlinear component of the CRRAM bid function that characterizes high bids, Kagel and Levin (1985) looked at individual subject bid functions, truncating the data set consistent with the requirements of CRRAM and testing whether CRRAM's restrictions held as expected profits conditional on winning the auction varied. Kagel and Levin (1985) reject the null hypothesis that the bid coefficients satisfy CRRAM's restrictions for 40 percent of their subjects.[30]

Further evidence that subjects increase their bids proportionately more as expected profit increases is reported in Kagel and Roth (1992), in this case for auctions with varying numbers of bidders. In these auctions, the underlying support from which valuations were drawn remained constant as n decreased, so that the expected profit conditional on winning the auction increased. Actual profit as a percentage of the predicted RNNE profit decreased in eight out of eight auction series. In contrast, according to CRRAM, profit earned as a percentage of RNNE profit should increase with decreases in n.[31]

These results stand in marked contrast to the tripling of payoff values reported by Cox et al. (1983b, 1988) which resulted in no significant differences in mean intercept values or slopes of individual subject bid functions and no significant changes in market prices. These differences in outcomes may reflect the fact that: (i) in Cox et al.'s experiment profits were tripled by increasing the conversion rate from experimental dollars to U.S. currency, so that the distribution of private values and all other variables over which bidders were choosing remained unchanged, requiring no changes in behavior to remain faithful to CRRAM, while (ii) the experiments reported in Kagel and Roth and Kagel et al. (1987) required substantial adjustments in bidding patterns if behavior was to satisfy CRRAM, thereby providing a much more demanding test of the null hypothesis. That is, Cox et al. (1983b, 1988) implemented an experimental manipulation for which their hypothesis predicted no change in subjects' behavior, and they detected no change and concluded that this supports their hypothesis. But if bidding above the RNNE were due to factors *not* accounted for by their CRRAM model, their experimental manipulation (which left unchanged virtually all of the experimental environment) was unlikely to change these factors either. In contrast, the manipulations reported in Kagel et al. (1987) and Kagel and Roth require a *change* in behavior to be faithful to CRRAM, and the predicted pattern of behavior was not observed.

Finally, Smith and Walker (1993) test for CRRA in a series of auctions with changes in conversion rates from experimental to U.S. dollars ranging between 0 and 20, and find that "all three regressions for risk aversion show that there is some tendency for individuals to bid higher (more risk averse) as payoff level is increased after correcting for the effect of experience" (242). Thus, at this stage, there is a growing consensus that, contrary to the CRRAM hypothesis, as expected profit conditional on winning the auction increases, subjects act as if they are relatively more risk averse.[32]

2. The Flat Maximum Critique

Harrison (1989) presents a methodological critique of the evidence Cox, Smith, and Walker employ to reject RNNE bidding theory. Harrison argues that under the typical payoff values employed the expected cost of deviations from the RNNE is quite small (less than \$0.05 at the median), so that in terms of expected monetary payoffs ("payoff space") many subjects had little to lose from deviating from the RNNE strategy (i.e., the payoff function around the maximum is flat). Harrison argues the significance of the differences reported between bids and the RNNE (deviations in the "message space") need to be reexamined.

In replying to Harrison, Cox et al. (1989a, 1989b, 1992) focus on Harrison's contention that "it is more natural to evaluate subject behavior in expected payoff space" (Harrison 1989, 749), and they interpret his cost of deviation measure as attaching cardinal value to expected utility differences. In arguing that it is more natural to evaluate behavior in payoff space, Harrison has overstated his case. However, as a diagnostic tool, evaluating outcomes in payoff space reflects the

common sense observation that when financial incentives are "small" (i.e., there are "low" expected costs to deviating from a particular outcome) the experimenter *may* have lost control, other arguments in subject's utility function *may* guide behavior, and this *may* result in systematic deviations from the theory. The suggestion is that bidding above the RNNE, and the differences in individual subject bid functions observed, *may* be characteristic of such systematic deviations.

Harrison's criticism applies with special force to Cox, Smith, and Walker's efforts to develop CRRAM since there is no known analytic solution for the CRRAM equilibrium bid function for bids above b^*, the maximum bid of the most risk loving subject in the auctions (Cox et al. 1988). Since the expected cost of deviating from the RNNE, or CRRAM, equilibrium bid function increases with bidder's private valuations (as the likelihood of winning the auction is higher for those with higher valuations), Cox, Smith, and Walker's investigation of CRRAM is based on private valuations for which the expected cost of deviating from equilibrium is the lowest. Kagel and Roth (1992) have examined this implication of Harrison's critique, computing the simple correlation coefficient between private values and the absolute value of the size of the deviation from RNNE bidding relative to the underlying private valuation ([actual bid – RNNE predicted bid]/private valuation). They report statistically significant negative correlations within each auction series, demonstrating that with lower valuations (and lower probability of winning the auction), the absolute size of the deviation from the RNNE is proportionately larger. Further, the average absolute size of these deviations is 2–3 times larger in cases where bids are made on the basis of the lowest 20 percent of the distribution of private values compared to the highest 20 percent of the distribution. That is, the greatest proportionate deviations from the RNNE predictions are made by bidders who draw the lowest valuations, bidders who have the smallest chance of winning the auction.[33] For these bidders the expected cost of deviating from the Nash equilibrium (even if they are risk neutral) is lowest.

The natural response to the flat maximum critique is to increase the marginal expected cost of deviations from the RNNE and see if they persist. Alternatively, one can implement experimental treatments designed to provide stern tests of the comparative static implications of the risk aversion hypothesis and see if it continues to organize behavior. Sections I.C.2 and I.D have already described two tests of the comparative static implications of risk aversion; additional tests are reported in sections I.G.3–I.G.4. The remainder of this section focuses on the effect of increasing the marginal cost of deviations from RNNE bidding.

Kagel and Roth (1992) report results from first-price auctions in which the expected profit conditional on winning the auction (assuming RNNE bidding) is substantially higher than in the experiments reported in Harrison (1989) or those typically conducted by Cox, Smith, and Walker. Looking at these experiments, median foregone expected income averages a little over \$0.18 in each auction period, with a high of \$0.31 and a low of \$0.05.[34] This is substantially higher than the median foregone expected income reported in Harrison (1989), which averages \$0.04, with a high of \$0.06 and a low of \$0.02. In addition, Smith and Walker (1993) have conducted auctions in which experimental dollars are transformed into U.S. dollars at rates of 10 and 20 to 1. The impact of this manipulation, evaluated in terms of its effect on individual subject bid functions, has been to reduce the mean square residual error as payoffs increase, so that subjects adhere closer to a given bidding strategy as payoffs increase, and for subjects to bid as if they are somewhat more risk averse. The median foregone expected income in auctions with the highest dollar transformations in Smith and Walker (1993) are substantially greater than those reported in Cox et al. (1988) or in the Harrison (1989) paper. Of course it is not possible to determine if these higher levels of foregone expected income, or any level of foregone expected income for that matter, are sufficiently large to be salient for subjects. But scaling up the expected loss function, to these levels at least, does not eliminate bidding above the RNNE.

In contrast to the relatively small median foregone expected income for all bidders, high value holders forego substantially more in expected income terms as a consequence of bidding above the RNNE, averaging \$1.34 for the auctions reviewed in Kagel and Roth (1992), compared to \$0.18 for all bidders. The difference in foregone expected income between the high private value holders and all bidders results from the fact that bidders with lower resale values have substantially smaller chances of winning the auction, which drives down foregone expected income substantially. The net effect is that bidders with relatively low resale values have sharply reduced financial incentives to behave in accordance with the theory. This is the primary point of Harrison's paper as I see it, a point which has been reflected in the data at several other places; for example, in the responses to varying numbers of bidders (section I.C) and the effects of public information in auctions with affiliated private values (section I.D).

In replying to his critics Harrison (1992) implicitly rejects the notion that increasing the expected cost of deviating from the RNNE and determining that the behavior persists (or gets stronger) is sufficient to overcome the flat maximum critique:

> An alternative class (of priors about decision making costs) that might be appropriate is the *percent* foregone income, with this *percentage* calculated relative to some reference decision. . . . Moreover, if subjects do follow a decision rule of this kind one could not eradicate a lack of payoff dominance by simply scaling up the payoff function (such procedures would work if we replaced "α %" with "α cents" in the decision rule). (Harrison 1992, 1441, italics added)

This idea of replacing an α cents rule with an α percent rule is simply not tenable for profit maximizing economic agents, the null hypothesis underlying Harrison's original argument and much of economic theory. Further, Harrison (1992) belittles the usefulness of comparative static tests of the internal consistency of proposed explanations for behavioral anomalies. This methodological position provides overly strong protection for the null hypothesis since many times economic models have inherently flat maxima. Therefore, tests of behavioral consistency across a variety of treatment conditions provide a basis for investigating such

anomalies. To dismiss such anomalies strictly because the null hypothesis has a relatively flat maximum simply affords the null hypothesis too much protection, particularly in the absence of any hard data that the theory works as predicted. (Tests of the internal consistency of the risk aversion hypothesis across private value auction environments are reported in the next section.)

3. Risk Aversion and Overbidding in Related Environments

Bidding above the RNNE in first-price auctions can be explained in terms of risk aversion. Empirically, in fitting bid functions to the data, what risk aversion does is add an extra degree of freedom. Accounting for diversity in bidder's risk preferences adds even more degrees of freedom, which is bound to result in an even closer fit to the data. The question remains though whether risk aversion and bidder heterogeneity provide a coherent explanation of deviations from RNNE bidding theory across a variety of auction institutions and experimental manipulations. To the extent that it does, there is added support for the argument that risk aversion is the mechanism underlying bidding above the RNNE in first-price auctions. To the extent that it does not, it is doubtful that risk aversion is the basic causal factor behind such overbidding. Note, in raising the issue of how well risk aversion does in related auction environments I am not denying the utility of exploring CRRAM (or any other model for that matter) in efforts to organize deviations from risk neutral bidding in first-price auctions. What I am arguing, however, is that risk aversion has implications outside the domain of first-price auctions, and it is incumbent on experimenters to explore these implications before feeling too comfortable with the risk aversion argument, no matter how good the fit to data from first-price auctions.

Sections I.C.2 (the case of numbers uncertainty) and I.D (the effects of public information in auctions with affiliated private values) provide two instances in which risk aversion serves to account, reasonably well, for comparative static responses to novel treatment conditions. Further, Chen and Plott (1991) investigate first-price auctions where valuations are drawn from nonuniform distributions and where the Nash equilibrium bidding function is nonlinear. Here too CRRAM exhibits good fits to the data. These results—in conjunction with the fact that bidding above the RNNE tends to be most pronounced in auctions with higher expected profit conditional on winning the auction, conditions under which risk aversion would be expected to be most prominent in the data—provide evidence that risk aversion has some explanatory power in private value auctions. However, there are several settings where bid patterns directly contradict the risk aversion hypothesis or raise questions about the role of risk aversion in bidding above the RNNE in first-price auctions.

The persistent bidding above the dominant strategy in second-price auctions (section II.A.2) can't help but raise the question of whether similar perceptual errors might be responsible for bidding above the RNNE in first-price auctions. Note that bidding the dominant strategy in second-priced auctions is *independent* of attitudes towards risk. Nevertheless, subjects persistently bid above their private valuations when they are permitted to do so. In light of this, it is natural to entertain the hypothesis that related perceptual errors may help explain systematic deviations from the RNNE in first-price auctions (see the discussion in section I.B.3c).[35] In fact Cox et al. (1985b) suggest one such model in a footnote, although they do not pursue it relative to the CRRAM alternative.

In Cox et al. (1984) multiple unit discriminatory auctions were conducted under several different values of n (the number of bidders) and q (the quantity offered for sale), with data being reported for experienced bidders (in discriminatory auctions the winning bidders each pay the price bid). In four of the ten treatment conditions studied, average revenue was consistently *lower* (usually significantly lower) than predicted under the RNNE bidding model, contrary to the risk aversion hypothesis. In the other six treatment conditions revenue was consistently greater (usually significantly greater) than predicted under RNNE, consistent with the risk aversion hypothesis.

Cox et al. (1984) explore a number of potential explanations for these differences. They are unable to find any explanation that works consistently across treatment conditions:

> But at this juncture in the research program we have no explanation of the Group I experiments in Figure 12 that are inconsistent with both the CRRA and the VHR models (models of risk averse Nash equilibrium bidding). The group I experiments are characterized by a pronounced tendency of individuals to bid below the risk neutral Vickrey bid function. It is natural to conjecture that this is due either to cooperative behavior as suggested in [Cox, Robertson, and Smith 1982] in single unit auctions when there are only three bidders, or to strictly convex (risk preferring) preferences for monetary outcome. (Cox, Smith, and Walker 1984, 1008).

Note that the cooperative behavior referred to in single unit auctions with three bidders has since been resolved (the results do not replicate with higher expected profits and/or somewhat different procedures—section I.B.1). The anomalous findings relative to risk aversion for multiple unit auctions have never been resolved.

In third-price IPV auctions risk aversion involves bidding *below* the RNNE prediction, the exact opposite of the pattern implied in first-price auctions (Kagel and Levin 1993). In auctions with five active bidders this prediction is supported, as the majority (61 percent) of bids lie below the RNNE prediction. However, in auctions with ten bidders, the majority of bids (60 percent) lie *above* the RNNE line. In going from five to ten bidders subjects first act as if they are predominantly risk averse, then as if they are predominantly risk loving! The suggestion is that something other than risk attitudes, as specified in conventional expected utility theory, is guiding behavior here.

Cason (1992) studied an auction in which several buyers bid for a single item whose price is fixed by a random draw from a known distribution. The value bidders place on the item is known at the time they bid, but the randomly determined sale price is unknown. The high bidder wins the item and pays the ran-

domly determined sale price, provided his bid is at or above it. In this set-up the symmetric RNNE calls for bidding above resale values (as bids only determine whether or not a buyer wins the item, not the price paid). With either constant absolute or constant relative, risk aversion equilibrium bids are *below* the symmetric RNNE (as in a third-price auction). Nevertheless Cason reports bidding significantly above the RNNE in auctions with both 3 and 6 bidders and with both inexperienced and experienced bidders.

Harrison (1990) attacks the issue of risk aversion directly. Using the Becker, DeGroot, Marschak (1964) procedure (see Roth, chapter 1, section III.F.4.a) to test directly for risk aversion over small gambles, he determines the distribution of risk preferences for a group of undergraduate students, not unlike those used in typical first-price auction experiments. Using Bayesian econometric techniques and data from Cox et al. (1982, 1983a, 1983b), he finds far too much risk loving in his sample population to be consistent with the degree of risk aversion implied by the CRRAM hypothesis using the first-price auction data.[36] Note that Harrison's conclusions are restricted to the CRRAM specification as it is not possible, in general, to recover risk preferences from bidding since this reflects bidders' own risk preferences *and* their reaction to the behavior (risk preferences) of others, unlike the gambles which strictly measure own risk preferences. CRRAM however has the special property that the linear portion of the bid function depends exclusively on bidders' own risk preferences, being independent of the distribution of risk preferences in the sample population. For more general specifications of risk aversion, consistency tests would require playing against programmed bidders who played according to the RNNE. To be most relevant tests of this sort should use reference gambles that are fully comparable to those faced in the auctions, along with using the same subjects in both the gambles and the auctions.

4. Using the Binary Lottery Procedure to Control for Risk Aversion

Assuming that subjects maximize expected utility, experimenters can, at least in theory, use a binary lottery procedure to induce risk preferences of their choosing (Roth and Malouf 1979, Berg, Daley, Dickhout, and O'Brien 1986). Under the lottery procedure, participants in the auction earn points (rather than income) which are used to determine their chances of winning the larger of two monetary prizes in a lottery following the auction (see the discussion in section III.F.4 of the Introduction).[37] Since according to expected utility theory preferences are linear in the probabilities and bidders' earnings are in probability points, as long as there is a one-to-one relationship between points and lottery tickets, the lottery procedure must lead expected utility maximizers to bid as if they are risk neutral.

There are at least three different ways of looking at efforts to employ the binary lottery procedure to control bidders' risk preferences in auctions. First, these efforts may be viewed as an additional comparative static test of the CRRAM hypothesis, and of risk aversion in general, since if bidders are indeed expected utility maximizers, then it follows as a matter of logic that the binary lottery procedure is capable of controlling risk preferences (Kagel and Roth 1992, express this point of view). Second, directly contrary to this point of view, auctions may be viewed as a vehicle for testing the binary lottery procedure, so that failure to control risk preferences in auctions would result in questioning the general effectiveness of the lottery procedure. This application involves a joint test of Nash equilibrium bidding theory and the lottery procedure, so that the validity of the test depends critically on one's priors regarding the ability of Nash equilibrium bidding theory to organize precisely behavior in private value auctions (Cox et al. [1985b] and Walker, Smith, and Cox [1990] express this point of view). Third, applications of the binary lottery procedure may be viewed on strictly practical grounds as efforts to develop another tool to help with theory testing; for example, if the binary lottery procedure can be used to move bidders closer to risk neutral bidding, more precise tests of the revenue raising effects of public information in auctions with affiliated private value can be conducted since increased revenue depends critically on the existence of risk neutral bidders (recall section I.D).

Cox et al. (1985b) and Walker et al. (1990) report results from a series of first-price private value auctions designed to test whether the lottery procedure does in fact induce risk neutral bidding.[38] The criteria for determining risk neutrality is whether the slope coefficient in the linear regression specification (equation (10)) is significantly different from the risk neutral prediction using ordinary least squares (OLS) regressions. Their primary conclusion is that the lottery procedure can *not* be used reliably to induce risk neutral bidding as the large majority of slope coefficients remain significantly above the RNNE prediction. There is some suggestion of prior conditioning effects in their data, as bidding under the lottery procedure moves subjects with no prior experience in first-price auctions in the direction of risk neutrality, compared to inexperienced subjects without the lottery procedure (although here too the average value of β is significantly above the risk neutral prediction; Walker et al. 1990, conclusion 6).[39]

Walker et al. (1990) rationalize the failure of the lottery procedure to induce risk neutral bidding and of their maintained hypothesis of CRRAM, which assumes expected utility maximization, by arguing that the failure of the lottery procedure is due to the failure of the compound lottery axiom.

> Our tests of the lottery procedure are based on the hypothesis that the CRRAM bidding model of first-price auctions is a successful predictor of bidding behavior. We chose the environment of the first-price auction because of the success of earlier direct tests of the CRRAM for that market institution. Given the lack of success of the lottery procedure in this environment, we feel our results have important implications for future research. First, it can be conjectured that our results provide indirect evidence of the failure of the compound lottery axiom of expected utility theory in the sense that there is an explicit compounding of the lottery and the auction procedure (in the tests of the compound lottery axiom). (23)

Walker et al. recognize that derivation of the CRRAM bidding function makes use of the compound lottery axiom so that they also argue

> In suggesting below that the failure of the lottery procedure to induce risk neutral behavior may be due to the failure of the compound lottery axiom we do not suggest that the axiom need fail in all contexts. In fact, within the context of CRRAM there is no evidence this axiom fails. (23, note 1)

Since Walker et al. provide no references to any independent experiments designed to isolate tests of the compound lottery axiom within the context of CRRAM, to this reviewer, at least, the results of Walker et al. might better be viewed as further evidence that risk aversion is not the sole, or even the primary, factor underlying bidding above the RNNE in first-price auctions (recall section I.G.3).

Rietz (1993) raises questions about the conclusions of Cox et al. (1985b) and Walker et al. from another perspective and reaches quite different conclusions regarding the ability of the lottery technique to control bidders' risk preferences. A potentially important procedural difference between Rietz and Walker et al. (1990), (as well as Cox et al. 1985a), is that Walker et al. always quoted values, bids, and profits in monetary units, which then serve to determine the probability of winning the large prize. In contrast, Rietz quotes all values, bids, and profits directly in terms of points. In addition, Rietz examines the effects of three different types of prior experience: (1) bidders start out in (standard) first-price dollar auctions, (2) the first-price lottery procedure is used throughout, and (3) the lottery procedure is first employed in second-price auctions, to familiarize subjects with the procedure when there is a "simple dominant strategy" (Rietz did not permit subjects to bid above their private value in the second-price auctions, simplifying the decision problem considerably).

Rietz finds that prior experience is important. Bidders who start out in standard first-price dollar auctions move closer to the RNNE prediction following the cross-over to the lottery procedure. But a null hypothesis of risk neutrality can still be rejected under the lottery technique.[40] In auction series employing the first-price lottery procedure from the start, the price distribution is even closer to the RNNE prediction. In this case the price distribution in one of two auction series does not differ significantly from the RNNE prediction. Finally, for bidders who start out in second-price lottery auctions, none differed from the RNNE prediction following the cross-over to first-price lottery auctions. With this kind of prior training the lottery procedure was successful in inducing the desired preferences. Unfortunately Rietz reports no control condition using second-price dollar auctions that were then crossed over to first-price dollar auctions to determine if the second-price training itself would produce closer conformity to the RNNE outcome. However, barring a hysteresis effect of this sort, which seems unlikely (see note 11 for evidence supporting this presumption), his results suggest that successful implementation of the lottery procedure rests on subjects first becoming familiar with the procedure in a fairly simple context and then placing them in more complicated settings like the first-price auction.

Rietz also looks at individual subject bid functions. He compares tests of the RNNE hypothesis using OLS estimators with least absolute deviation (LAD) estimators (Powell 1984). Rietz argues that censoring (it is individually irrational for subjects to bid above their resale values in first-price auctions) and heteroskedastic errors result in artificially high rejection rates of the null hypothesis of risk neutrality using OLS estimators. The LAD estimators correct these biases. The differences can be quite dramatic. Using OLS estimators, Rietz rejects the null hypothesis that the binary lottery procedure controlled subjects' preferences, as intended, for seven of eight subjects in series that began with dollar auctions, for four of eight subjects in series that always employed first-price lottery auctions, and for seven of sixteen subjects in auctions that began with the second-price lottery procedure. The corresponding rejection rates using the LAD estimators are 1 of 8, 0 of 8 and 0 of 16. Although Rietz's LAD tests are based on asymptotic properties of the estimators rather than their small sample properties, which may understate the rejection rate, he argues that the results indicate that OLS estimates, like those employed in Cox et al. (1985b) and Walker et al. (1990), may overstate the rejection rate. The econometric biases Rietz identifies in estimating individual subject bid functions are likely to apply to other tests of auction theory using OLS estimates as well.

5. Summing Up

This section started by noting that there was considerable debate among experimenters regarding the role of risk aversion in bidding above the RNNE outcome in first-price private value auctions. The debate has pushed research in a number of directions. Cox, Smith, and Walker's initial work demonstrated that a model of heterogeneous risk averse bidders provided a better fit to the experimental data than homogeneous risk neutral, or risk averse, bidding models. At the same time CRRAM has some notable inadequacies even in first-price private value auctions: there are (mild) nonlinearities in individual subject bid functions, intercepts of individual subject bid functions are often significantly different from the CRRAM prediction of zero, and in a number of cases bidders act as if they become more risk averse as the expected profit conditional on winning the auction increases. There is some suggestion that minimum income thresholds from winning the auction may be responsible for these last two effects.

The flat maximum critique argues that deviations from risk neutrality observed in first-price private value auctions may be a consequence of low expected costs of deviation from RNNE, so that the resulting loss of experimental control permits other arguments in subject's utility function to guide behavior. While a number of experimenters have taken the flat maximum critique as a frontal attack on auction experiments (and in some cases on the entire field of experimental economics), to this reviewer it serves as a useful diagnostic tool that *may* indicate when control of subjects' behavior is likely to be lost. Results from a number of first-price auction experiments involving considerably greater expected cost of deviating from the RNNE than in Harrison's study or the typical Cox, Smith, and

Walker experiment show subjects continuing to bid above the RNNE at considerably greater expected cost to themselves. Thus, scaling up the expected loss function, to these levels at least, does not eliminate bidding above the RNNE.

Looking at bidding in other auction settings, there are several instances in which the risk aversion hypothesis fails, most notably in third-price auctions and in multiple unit discriminatory auctions. These results, in conjunction with bidding above the dominant strategy in second-price auctions, raise serious questions regarding the role of risk aversion as the mechanism underlying bidding above the RNNE in single unit first-price auctions. One response to these contrary data is to point out that they come from considerably more complicated settings than first-price auctions, so that to some extent these breakdowns are not surprising and should not be counted too heavily against the risk aversion explanation.[41] To this reviewer, this response is inadequate since it fails to account for the fact that the very simplicity of first-price auctions means that bidders can develop reasonable ad hoc rules of thumb that result in "as if" risk averse Nash equilibrium bidding, when in fact the underlying behavioral process bears little relationship to Nash equilibrium bidding theory (see equation (5) in section I.C.1). Consequently, it requires examining behavior in more complicated environments to distinguish between risk aversion and such ad hoc bidding rules. The bottom line to all of this is that despite its failures, the risk aversion explanation has its successes too, so that it probably has some role to play in explaining bidding in private value auctions.

Applications of the binary lottery procedure to induce risk neutral bidding in first-price auctions has met with mixed results. Rietz (1993) reports considerably more success than Cox et al. (1985b) or Walker et al. (1990) in inducing risk neutral preferences as judged by the equilibrium predictions for risk neutral bidders. There are differences between the two experiments both in terms of underlying procedures and statistical tests used to evaluate the success of the technique.

II. Common Value Auctions

In common value auctions the value of the auctioned item is the same to all bidders. What makes the auction interesting is that bidders do not know the value at the time they bid. Instead they receive signal values that are related to (affiliated with) the value of the item.[42] Mineral lease auctions, particularly the Federal government's outer continental shelf (OCS) oil lease auctions, are common value auctions. There is a common value element to most auctions. Bidders for an oil painting may purchase for their own pleasure, a private value element, but they may also bid for investment and eventual resale, reflecting the common value element.

Judgmental failures in common value auctions are known as the "winner's curse." Although all bidders obtain unbiased estimates of the item's value, assuming homogeneous bid functions, they only win in cases where they have the highest signal value. Unless this adverse selection problem is accounted for in the bidding, it will result in winning bids that produce below normal or even negative profits. The systematic failure to account for this adverse selection problem is referred to as the "winner's curse."

Oil companies claim they fell prey to the winner's curse in early OCS lease sales (Capen, Clapp, and Campbell 1971; Lorenz and Dougherty 1983, and references cited there-in). Similar claims have been made in auctions for book publication rights (Dessauer 1981), professional baseball's free agency market (Cassing and Douglas 1980), corporate takeover battles (Roll 1986), and real estate (Ashenfelter and Genesore 1992). Economists typically treat such claims with caution as they imply that bidders repeatedly err, in violation of basic notions of economic rationality (see the exchange between Cox and Isaac [1984, 1986] and Brown [1986], for example). Further, a number of studies question the oil companies claims regarding losses in early OCS sales (see Wilson [1992] for a brief review of the literature; see also section III.B below).

Laboratory experiments show that inexperienced bidders are quite susceptible to the winner's curse (Bazerman and Samuelson 1983; Kagel and Levin 1986; Kagel, et al. 1989; reviewed in Roth 1988; chapter 1, section III.E of this volume). In fact, the winner's curse has been such a pervasive phenomenon that most of the initial experimentation in the area has focused on its robustness and the features of the environment that might attenuate its effects. Secondary issues concern the role of public information and different auction institutions on revenue. Public information concerning the value of the item should raise prices, assuming risk neutral bidders and no winner's curse (Milgrom and Weber 1982). Further, second-price auctions should raise more revenue than first-price auctions, and English auctions should raise more revenue than second-price auctions, assuming risk neutral bidders, no winner's curse, and a symmetric Nash equilibrium (Milgrom and Weber 1982).

Section II.A looks at the winner's curse in sealed bid common value auctions. Experimental procedures are first characterized, followed by a summary of results in first-price auctions, including the effects of public information on revenue. The impact of limited-liability for losses is discussed and results are reported from second-price auctions. Section II.B looks at the winner's curse in English auctions and in auctions with asymmetric information, where one bidder knows the value of the item with certainty. Experimental work on the winner's curse in other market settings—in bilateral bargaining games with uncertainty, in "blind bid" auctions, and in two-sided auction markets with a lemon's problem—are reviewed in section III.C. What little is known about whether and how bidders learn to overcome the winner's curse is reported in section II.D.

A. *The Winner's Curse in Sealed Bid Common Value Auctions*

Procedures employed in common value auction experiments are similar to auctions with affiliated private values (section I.D). The common value, x_o, is chosen randomly each period from a uniform distribution on $(\underline{x}, \bar{x})$. Under symmetric information conditions each bidder is given a private information signal, x, drawn

from a uniform distribution on $[x_o - \epsilon, x_o + \epsilon]$, where ϵ is known (thus private information signals are positively affiliated). In first-price auctions the high bidder earns $x_o - b$, where b is the high bid. Losing bidders earn zero profits.

Under this design bidders have a risk free strategy of bidding max $[x - \epsilon, \underline{x}]$, which is computed for them in the form of a lower bound estimate of x_o, along with an upper bound estimate of x_o (min $[x + \epsilon, \bar{x}]$). Bidders are provided with illustrative distributions of signal values relative to x_o and several dry runs before playing for cash. Following each auction period they are provided with the complete set of bids, listed from highest to lowest, along with the corresponding signal values, the value of x_o and the earnings of the high bidder (subject identification numbers are, however, suppressed).

Bidders are provided with a starting cash balance and have the opportunity to bid in a series of auctions. Should a subject exhaust her cash balance, she is declared bankrupt and no longer allowed to bid. Surviving bidders are paid their end of experiment balances in cash. To hold the number of bidders fixed while controlling for bankruptcies, $m > n$ subjects are often recruited, with only n bidding at any given time (who bids in each period is determined randomly or by a fixed rotation rule). As bankruptcies occur m shrinks, but (hopefully) remains greater than or equal to the target value n. Alternative solutions to the bankruptcy problem are discussed below.

1. First-Price Sealed Bid Auctions

Wilson (1977) was the first to develop a Nash equilibrium solution for first-price common value auctions, while Milgrom and Weber (1982) provide some significant extensions and generalizations of the Wilson model. For risk neutral bidders and signals in the interval $\underline{x} + \epsilon \leq x \leq \bar{x} - \epsilon$, the Nash equilibrium bid function is

$$b(x) = x - \epsilon + Y \tag{9}$$

where Y is a negative exponential, which becomes negligible rapidly as x moves beyond $\underline{x} + \epsilon$.[43]

In common value auctions bidders usually win the item when they have the highest, or one of the highest, estimates of value. Define $E[x_o \mid X = x_1]$ to be the expected value of the item conditional on having the highest signal value.[44] The latter provides a convenient measure of the extent to which bidders suffer from the winner's curse since in auctions in which the high signal holder always wins the item, bidding above $E[x_o \mid X = x_1]$ results in negative expected profit. Further, even with zero correlation between bids and signal values, if everyone else bids above $E[x_o \mid X = x_1]$, bidding above $E[x_o \mid X = x_1]$ results in negative expected profit as well. As such, if the high signal holder frequently wins the auction and a reasonably large number of rivals are bidding above $E[x_o \mid X = x_1]$, individuals bidding above $E[x_o \mid X = x_1]$ are likely to earn negative expected profit.

Auctions with inexperienced bidders show a pervasive winner's curse that results in numerous bankruptcies. Table 7.5 provides illustrative data on this point. For the first nine auction periods profits averaged −\$2.57 compared to the RNNE

Table 7.5.
Profits and Bidding in First Nine Auction Periods for Inexperienced Bidders

Experiment	Percentage of Auctions with Positive Profits	Average Actual Profits (t-statistic)	Average Predicted Profits under RNNE (S_M)[a]	Percentage of All Bids with $b > E[x_0 \mid X = x_1]$	Percentage of Auctions Won by High Signal Holder	Percentage of High Bids with $b_1 > E[x_0 \mid X = x_1]$	Percentage of Subjects Going Bankrupt[b]
1	0.0	−4.83 (−3.62)[c]	0.72 (0.21)	63.4	55.6	100.0	50.0
2	33.3	−2.19 (−1.66)	2.18 (1.02)	51.9	33.3	88.9	16.7
3	11.1	−6.57 (−2.80)[d]	1.12 (1.19)	74.6	44.4	88.9	62.5
4	11.1	−2.26 (−3.04)[c]	0.85 (0.43)	41.8	55.6	55.6	16.7
5	33.3	−0.84 (−1.00)	3.60 (1.29)	48.1	44.4	88.9	50.0
6	22.2	−2.65 (−1.53)	2.55 (1.17)	67.3	66.7	100.0	33.3
7	11.1	−2.04 (−2.75)[d]	0.57 (0.25)	58.5	88.9	66.7	50.0
8	11.1	−1.40 (−2.43)[d]	1.59 (0.34)	51.9	55.6	55.6	16.7
9	44.4	0.32 (0.30)	2.37 (0.76)	35.2	88.6	66.7	16.7
10	0.0	−2.78 (−3.65)[c]	3.53 (0.74)	77.2	66.7	100.0	20.0
11	11.1	−3.05 (−3.53)[c]	1.82 (0.29)	81.5	55.6	88.9	37.5
Average	17.2	−2.57	1.90	59.4	59.6	81.8	41.1

Source: Kagel, Levin, Battalio, and Meyer 1989.
[a] S_M = standard error of mean.
[b] For all auction periods.
[c] Statistically significant at the 1 percent level, two-tailed test.
[d] Statistically significant at the 5 percent level, two-tailed test.

prediction of \$1.90, with only 17 percent of all auction periods having positive profits. This is not a simple matter of bad luck either, as 59 percent of all bids and 82 percent of the high bids were above $E[x_o \mid X = x_1]$. Further, 40 percent of all subjects starting these auctions went bankrupt. The winner's curse for inexperienced bidders is a genuinely pervasive problem which has been reported under a variety of treatment conditions (Kagel et al., 1989, Lind and Plott 1991) and for different subject populations, including professional bidders from the commercial construction industry (Dyer, Kagel, and Levin 1989b, discussed in section III.B below).[45]

Kagel and Levin (1986) report auctions for moderately experienced bidders (those who had participated in at least one prior first-price common value auction experiment). Treatment variables of interest were the number of rival bidders and the effects of public information about x_o on revenue. Table 7.6 reports some of their results. For small groups (auctions with three or four bidders), the general pattern was one of positive average profits which, although well below the RNNE prediction, were clearly closer to the RNNE than the zero/negative profit levels of the winner's curse (profits averaged \$4.68 per period, about 65 percent of the RNNE prediction). In contrast, larger groups (auctions with six or seven bidders) had average profits of –\$.88, compared to the RNNE prediction of \$4.65 per period. Although profits earned were substantially better than predicted under a naive bidding model, indicating considerable adjustment to the adverse-selection problem, these adjustments were far from complete.[46] Further, comparing large and small group auctions, actual profit decreased substantially more than profit opportunities as measured by the RNNE criteria. This implies that subjects were bidding more aggressively, rather than less aggressively, as the number of rivals increased, contrary to the RNNE prediction.

Public information was provided to bidders in the form of announcing x_L, the lowest signal value (and announcing that x_L was announced). For the RNNE, public information about the value of the item raises expected revenue. The mechanism here is similar to the one operating in auctions with affiliated private values: Each bidder bids as if she has the highest signal value since this is when she wins, the only event that counts. For bidders whose private information signals are less than x_1, public information about the value of the item will, ex post, raise the expected value of the item. This results in an upward revision of these bids, which in turn puts pressure on the highest signal holder to bid more.

These strategic considerations hold for a wide variety of public information signals (Milgrom and Weber 1982). There are, however, several methodological advantages to using x_L. First, the RNNE bid function is readily solved for x_L, so that the experimenter continues to have a benchmark model of fully rational behavior against which to compare actual bidding. Second, x_L provides a substantial amount of public information about x_o (it cuts expected profit in half), while still maintaining an interesting auction. As such it should have a substantial impact on prices, regardless of any inherent noise in behavior. And the experimenter can always implement finer, more subtle probes of public information after seeing what happens with such a strong treatment effect.

Table 7.6. Profits and Bidding by Experiment and Number of Active Bidders: Private Information Conditions (profits measured in dollars)

Auction Series (Number of Periods)	Number of Active Bidders	Average Actual Profit (t-Statistics[a])	Average Profit under RNNE (Standard Error of Mean)	Percentage of Auctions Won by High Signal Holder	Percentage of High Bids with $b_1 > E[x_0 \mid X = x_1]$
6 (31)	3–4	3.73 (2.70)[b]	9.51 (1.70)	67.7	22.6
2 (18)	4	4.61 (4.35)[c]	4.99 (1.03)	88.9	0.0
3 small (14)	4	7.53 (2.07)	6.51 (2.65)	78.6	14.3
7 small (19)	4	5.83 (3.35)[c]	8.56 (2.07)	63.2	10.5
8 small (23)	4	1.70 (1.56)	6.38 (1.21)	84.6	39.1
1 (18)	5	2.89 (3.14)[c]	5.19 (0.86)	72.2	27.8
3 large (11)	5–7	–2.92 (–1.49)	3.65 (0.62)	81.8	63.6
7 large (18)	6	1.89 (1.67)	4.70 (1.03)	72.2	22.2
4 (25)	6–7	–0.23 (–0.15)	4.78 (0.92)	69.2	46.2
5 (26)	7	–0.41 (–0.44)	5.25 (1.03)	42.3	65.4
8 large (14)	7	–2.74 (–2.04)	5.03 (1.40)	78.6	71.4

Source: Kagel and Levin 1986.
[a] Tests null hypothesis that mean is different from 0.0.
[b] Significant at 5 percent level, two-tailed t-test.
[c] Significant at 1 percent level, two-tailed t-test.

Kagel and Levin (1986) found that in auctions with small numbers of bidders (three or four), public information resulted in statistically significant increases in revenue that averaged 38 percent of the RNNE model's prediction. However, in auctions with larger numbers of bidders (six or seven), public information significantly *reduced* average revenue. Kagel and Levin attribute this reduction in revenue to the presence of a relatively strong winner's curse in auctions with large numbers of bidders. If bidders suffer from a winner's curse, the high bidder consistently overestimates the item's value, and announcing x_L is likely to result in a downward revision of the most optimistic bidder's estimate. This introduces a potentially powerful offset to any strategic forces tending to raise bids and will result in reduced revenue if the winner's curse is strong enough. Kagel and Levin relate this result to anomalous findings from OCS auctions with field data (this relationship is summarized in Roth 1988; Introduction, section III.E and section III.B below).

Finally, looking at market outcomes with x_L announced, average profits were positive in all auction sessions, with no systematic differences in realized profits relative to predictions between auctions with small and large numbers of bidders. Further, although there was considerable variation in profits relative to the RNNE model's predictions across auction series, on average profits were only slightly less than predicted. These two characteristics suggest that with the large dose of public information involved in announcing x_L, the winner's curse had largely been eliminated.

2. Limited Liability and "Safe Havens"

In the Kagel and Levin (1986) design subjects enjoyed limited liability as they could not lose more than their starting cash balances. Hansen and Lott (1991) have argued that the overly aggressive bidding reported in Kagel and Levin *may* not have resulted from the winner's curse, but could instead be a rational response to this limited liability. In a single-shot auction, if a bidder's cash balance is zero, so that they are not liable for losses, it indeed pays to overbid relative to the Nash equilibrium. With downside losses eliminated the only constraint on more aggressive bidding is the opportunity cost of bidding more than is necessary to win the item. In exchange, higher bids increase the probability of winning the item and making positive profits. The net effect, in the case of zero or very small cash balances, is an incentive to bid more than the Nash equilibrium prediction.

Hansen and Lott's argument provides a possible alternative explanation to the overly aggressive bidding reported in Kagel and Levin and Kagel et al. (1989). It also serves as a warning to experimenters not to overlook potential limited-liability problems in designing their experiments. For example, research on bubbles in asset markets indicates that limited liability for losses may be at least partially responsible for the size and persistence of the bubbles reported (Bronfman 1990).

Responses to the limited-liability argument in common value auctions have been twofold. First, Kagel and Levin (1991) have revaluated their design and their data in light of Hansen and Lott's arguments. They show that their design protects against limited-liability problems and that for almost all bidders cash balances were *always* large enough that it *never* paid to deviate from the Nash equilibrium bidding strategy in a single-shot auction. Second, Lind and Plott (1991) replicated Kagel and Levin's experiment in a design that eliminated limited-liability problems and reproduced Kagel and Levin's (1986) primary results. This provides experimental verification that limited-liability forces do not account for the overly aggressive bidding reported.

Kagel and Levin's design protects against limited liability for losses since bidding $x - \epsilon$ insures against all losses and is close to the predicted RNNE bid function (recall equation (9)). For example, consider a bidder with a private information signal of \$80 in an auction where the value of the item is \$50, the bidder has a cash balance of \$10, and $\epsilon = \$30$. In this example the RNNE bid is \$52.27 in a market with four bidders, or \$50.41 in a market with seven bidders, so that the bidder is fully liable for all losses (and a good deal more) relative to the Nash equilibrium bid.[47] And as long as bidders have sufficient cash balances to cover their maximum losses relative to the Nash equilibrium bid, overbidding *cannot* be rationalized in terms of a Nash equilibrium.[48] Rather, overbidding in this case must be explained on some other grounds, such as the judgmental error underlying the winner's curse.[49]

Lind and Plott (1991) replicated Kagel and Levin's (1986) results for buyers' and sellers' auctions in which bankruptcy problems were eliminated. To get around the bankruptcy problem in buyers' markets, private value auctions where subjects were sure to make money were conducted simultaneously, thereby providing a source of funds which reduced the likelihood of bankruptcy in the common value auction. In addition, subjects agreed that if they wound up with losses they would work them off doing work-study type duties (photocopying, running departmental errands, etc.) at the prevailing market wage rate.[50] In sellers' markets, bidders tendered offers to sell an item of unknown value. Each bidder was given one item with the option to keep it and collect its value or to sell it. In this auction, all subjects earned positive profits, including the winner, but the winner could suffer an opportunity cost by selling the item for less than its true value.[51]

Lind and Plott's results largely confirm those reported by Kagel and Levin and their associates. First, a winner's curse exists, and although the magnitude and frequency of losses decline with experience, it persists (see Table 7.7). Second, the winner's curse does not result from a few "irrational" bidders, but almost all agents experience the curse and bid consistent with curse behavior. Finally, Lind and Plott test between alternative models of bidder behavior—comparing the RNNE bidding model with the naive bidding model offered in Kagel and Levin (1986). Since these models imply different sets of parameter restrictions on a common functional form, Lind and Plott compute F-statistics comparing the sum of squared errors of the unrestricted model with the restricted model, using the F-statistic as a measure of the relative goodness of fit of the competing models.

Table 7.7. Frequency of Losses for Winners in All Experiments

	Experiment				
	1. Buyers	2. Buyers	3. Sellers	4. Sellers	5. Sellers
Periods 1–10					
Number of periods of loss	8/10	8/10	5/10	6/10	5/10
Average profit per period	−7.90	−8.31	−0.075	−0.048	0.001
	(−7.90)	(−8.31)	(−29.80)	(−48.20)	(1.10)
Average RNNE profit per period	4.53	5.70	0.177	0.060	0.048
	(4.53)	(5.70)	(70.96)	(60.44)	(68.71)
Periods 11–20					
Number of periods of loss	4/10	2/7	3/7	2/10	7/10
Average profit per period	4.57	3.12	0.053	0.032	−0.016
	(4.57)	(3.12)	(21.00)	(31.60)	(−22.40)
Average RNNE profit per period	18.47	13.58	0.212	0.048	0.037
	(18.47)	(13.58)	(84.85)	(48.15)	(52.68)
Periods 21–30					
Number of periods of loss				3/10	5/10
Average profit per period				0.058	−0.004
				(58.40)	(−6.10)
Average RNNE profit per period				0.104	0.090
				(104.02)	(128.91)
Periods 31–40					
Number of periods of loss				2/5	8/10
Average profit per period				0.063	−0.033
				(62.80)	(−46.80)
Average RNNE profit per period				0.065	0.024
				(65.34)	(33.72)

Source: Lind and Plott 1991.

Notes: Profits are measured in dollars and francs; numbers in parentheses are francs. $N = 7$ in all auctions; ϵ = \$30 and 200 francs in buyers and sellers auctions respectively. The RNNE was calculated only in cases where $\underline{x} + \epsilon \leq x_i \leq \bar{x} - \epsilon$. Some of the winners' signals were not in this range, in which case no prediction was made for the RNNE.

They find that neither model organizes the data, but that the RNNE provides a better fit. This last result, in conjunction with the negative average profits reported, indicate that there was partial but incomplete adjustment to the adverse selection forces in Lind and Plott's auctions.

Cox and Smith (1992) have conducted common value auctions with a "safe haven"—bidders have a choice between participating in the auction or obtaining a certain payoff from a "safe haven." The safe haven payoff is an independent

private value draw from a uniform distribution with each potential bidder learning their safe haven value before deciding whether or not to participate in the auction. One idea behind the safe haven is that it may alleviate the winner's curse through eliminating experimenter demand effects: Without the safe haven, bidders who have not learned to bid profitably may continue to make suboptimal decisions in response to implicit experimenter demands that they continue to bid and try to win the item. With a safe haven this implicit demand to win the item is eliminated, as the experimenter has provided subjects with an alternative, sanctioned activity.

To test for a safe haven effect, Cox and Smith conduct auctions with a zero income safe haven and with a positive income safe haven employing Kagel and Levin's (1986) design. They find substantially more bankruptcies and more winner's curse (bidding above $E[x_o \mid X = x_1]$) with the zero income safe haven. They interpret these results as "providing strong support for the importance of including a positive-income safe haven in common value bidding theory and experiments. Without the positive-income safe haven, the winner's curse is rampant and the theory fails dramatically to predict the behavior of inexperienced subjects" (Cox and Smith, 1992, 26).

There are a number of problems with this conclusion. First, with respect to the question of whether demand induced effects underlie the winner's curse, Kagel and Levin's design offers a considerably better option than the zero income safe haven: bid $x - \epsilon$, a sanctioned activity that provides bidders with a positive probability of winning the auction and completely protects them from losses. That is, the Kagel and Levin design has built into it a "safe haven" option that is both sanctioned and offers considerably higher expected earnings than the zero income safe haven. Second, to the extent that bidders do not have the wherewithal to identify this safe haven, as Cox and Smith implicitly assume, explicit provision of a zero income safe haven should be sufficient to eliminate any experimenter demand effects. Nevertheless, the winner's curse is rampant in Cox and Smith's zero-income safe haven sessions, producing an average bankruptcy rate of 75 percent for inexperienced bidders, compared to an average bankruptcy rate of 41 percent for inexperienced bidders in the no safe haven experiment in Kagel et al. (1989).[52] Third, differences in bankruptcy rates and the size of the winner's curse between positive income and zero income safe havens fail to control for large differences in the number of active bidders between the two treatments (for example, in Cox and Smith's Group II sessions with inexperienced bidders, there are 3.4 active bidders on average with the positive income safe haven versus 5.6 bidders in the zero income treatment). It is well established, both in theory and experimentally, that in auctions with fewer bidders the adverse selection problem is less severe, hence less room for bankruptcies and bidding above $E[x_o \mid X = x_1]$. (Indeed, under the naive bidding model offered in Kagel and Levin [1986] and Lind and Plott [1991], expected profits are positive as long as there are three or fewer bidders). As to the broader question this study implicitly poses, whether naturally occurring auction markets adjust to the winner's curse by causing bidders

to exit the market, or whether some other adjustment mechanism(s) are at work (like the ones characterized in sections II.C.2, II.C.3, and III.B.2 below), this remains an open empirical question the resolution of which will require careful field studies of auction markets.

3. Second-Price Sealed Bid Auctions

Lind and Plott (1991) are puzzled by their finding that although there is a winner's curse, the RNNE model provides the best fit to the data: "A major puzzle remains: of the models studied, the best is the risk-neutral Nash-equilibrium model, but that model predicts that the curse will not exist" (344). They go on to comment, "Part of the difficulty with further study stems from the lack of theory about (first-price) common value auctions with risk aversion. . . . If the effect of risk aversion is to raise the bidding function as it does in private (value) auctions, then risk-aversion . . . might resolve the puzzle; but, of course, this remains only a conjecture" (344). Second-price auctions provide an ideal vehicle for exploring these conjectures.

In contrast to first-price auctions, behavior of risk averse bidders is well understood in second-price auctions with both *symmetric* risk averse bidders and with *asymmetric* risk averse bidders.[53] Overly aggressive bidding in second-price common value auctions cannot be rationalized in terms of bidders' risk aversion. Rather, best responses involve less aggressive bids for risk averse than for risk neutral bidders. Further, this comparative static implication of the theory holds for both risk averse symmetric and asymmetric Nash equilibria and even extends to auctions where the strategy profile is not an equilibrium (corresponding predictions in first-price auctions require symmetry and are conditional on risk attitudes and the underlying distribution of information at bidders' disposal).[54]

Kagel, Levin, and Harstad (1994) investigate these comparative static predictions along with the effects of public information in second-price common value auctions. Using a fixed effect regression model, and comparing auctions with four and five bidders to auctions with six and seven bidders, they find no response to increasing numbers of rivals for moderately experienced bidders. This directly contradicts the Nash equilibrium prediction. However, it is consistent with a naive bidding model in which bidders fail to account for the adverse selection problems inherent in winning the auction.

In auctions with four or five bidders public information in the form of announcing x_L raises average revenue some 16 percent of the symmetric RNNE model's prediction (but the increase is not significant at conventional levels). In contrast, in auctions with 6 or 7 bidders announcing x_L *reduces* average revenue by \$4.00 per auction period (which is significant at conventional levels), compared to the predicted increase of \$1.80 per period under the symmetric RNNE. As in the first-price auctions, the ability of public information to increase revenue appears to be conditional on eliminating the worst effects of the winner's curse as bidders earned positive average profits in private information auctions with four or five bidders and substantial negative profits in auctions with six or seven bidders.

4. Summing Up

A strong winner's curse is reported for inexperienced bidders in sealed bid common value auctions as high bidders earn negative average profits and consistently bid above the expected value of the item conditional on having the high signal value. Similar results are reported for both student subjects and professional bidders from the construction industry (Dyer et al. 1989b). Arguments that these results can be accounted for on the basis of bidders' limited-liability for losses have been shown to be incorrect (Kagel and Levin 1991; Lind and Plott 1991). Safe haven treatments have yet to identify any experimenter demand effects underlying the winner's curse.

In the absence of a winner's curse, public information tends to raise revenue as the theory predicts. However, with a winner's curse, public information reduces revenue as the additional information helps bidders to correct for overly optimistic estimates of the item's worth. These results are found in both first and second-price auctions. Finally, increased numbers of bidders produces no change in bidding in second-price auctions, contrary to the robust Nash equilibrium prediction that bids will decrease.

B. More Winner's Curse: English Auctions and First-Price Auctions with Asymmetric Information

Dan Levin and I have investigated the winner's curse in two other common value auction settings—English auctions and first-price auctions with asymmetric information—in efforts to identify conditions where it might be eliminated for inexperienced bidders. In both settings the winner's curse is alive and well, although it is clearly less severe in English than in first-price auctions.

1. English Auctions

In a symmetric RNNE of an English auction, the bidder with the low signal value (x_L) drops out of the auction once the price reaches his signal value.[55] The intuition here is roughly as follows: Given symmetry, the low signal holder knows that those remaining in the auction have higher signal values. But the low signal holder can't profit from this additional information since it is only revealed once price is greater than these remaining signal values; that is, price is already greater than the expected value of the item to the low signal holder.

The price at which the low bidder drops out of the auction reveals his signal value to the remaining bidders. Since with symmetry you only win when you have the high signal value, given the uniform distribution of signal values around x_o, $(x_L + x)/2$ provides a sufficient statistic for x_o (where x is the bidder's own signal). This sufficient statistic is the equilibrium bid in the symmetric RNNE, with the high bidder paying the price at which the next-to-last bidder dropped out. Expected profits in the English auction are roughly half the level predicted in first-price auctions (as long as $n > 2$).

Table 7.8. English versus First-Price Auctions: Inexperienced Bidders

	ε = 6		ε = 12	
	Actual Profit	Predicted Profit	Actual Profit	Predicted Profit
English auction	–1.87 (0.51)	0.89 (0.29)	–1.80 (0.77)	1.68 (0.40)
First-price auction	–3.85 (0.71)	0.99 (0.19)	–3.75 (0.89)	2.76 (0.53)
Difference: first-price less English	–1.98 (0.87)	0.10 (0.34)	–1.95 (1.19)	1.08 (0.65)

Source: Kagel and Levin 1992.
Notes: $n = 7$. Actual and predicted profits are mean values, with standard error of the mean in parentheses.

The information dissemination process in English auctions has much in common with first and second-price auctions when the lowest private information signal is announced—the primary difference is that in the English auction information revelation is endogenous rather than exogenous. In sealed bid auctions with public information, bidders consistently earned positive average profits, even though these same bidders lost money under private information conditions (Kagel and Levin 1986; Kagel et al. 1992). This suggests the possibility that the winner's curse will be attenuated or even eliminated in English auctions, even with inexperienced bidders.

However, this is not the case (Kagel and Levin 1992). Table 7.8 reports actual profit and predicted profit under the symmetric RNNE from a series of English auctions with inexperienced bidders, along with data from a series of first-price auctions using the same subject population. Bidders in the English auctions failed to avoid the winner's curse, earning profits which were, on average, significantly below zero. However, profits were higher in the English auctions compared to profits for inexperienced bidders in first-price auctions who suffered from an even stronger winner's curse.[56] Higher profits translate into lower revenue, so these results directly contradict the symmetric RNNE model's prediction that English auctions will raise more revenue than first-price auctions. Lower revenues from "public" information (other bidders drop out prices) is, however, consistent with the effect of announcing the lowest private information signal in sealed-bid auctions when bidders suffer from a winner's curse (Kagel and Levin 1986; Kagel et al. 1992; see section II.A above). However, unlike the sealed bid auctions, in the English auctions this correction effect of "public" information does not completely eliminate the winner's curse as bidders continue to earn negative average profits.

Inexperienced bidders are unable to avoid the winner's curse in the English auctions for two related reasons. First, contrary to the theory's prediction, the low signal holder does not drop out of the bidding when the price reaches his signal value. Rather, the lowest drop-out price is typically above the lowest signal value, averaging $2.69 per auction period for the data in Table 7.6 (with a standard error of the mean of 1.00). Second, high bidders fail to compensate for the overly optimistic estimate of x_0 inherent in these drop-out prices. In fact they exacerbate it somewhat, as the average difference between the high bid and the symmetric RNNE bid was $3.30 per auction period. The level of overoptimism involved in the low drop-out price and the level of overbidding relative to the symmetric RNNE prediction are positively correlated (r = .63), indicating that the more optimistic the estimate of the item's value inherent in the low drop-out price, the higher the winning bid. This suggests that announcing the low signal value in English auctions would further attenuate the winner's curse, for which there is some limited experimental evidence (Kagel, unpublished data).

2. Auctions with Asymmetric Information

In auctions with asymmetric information some bidders are better informed than others about x_0. Dan Levin and I have implemented an extreme form of asymmetric information where one bidder (the insider), chosen at random in each auction period, knows the value of the item with certainty, and this is common knowledge. All remaining bidders (the outsiders) receive private information signals as in the symmetric information design. The bidder with inside information faces a pure strategic problem in the sense that the lower he bids the greater his profits will be should he win the item, but the less likely he is to win assuming that the others enter minimally competitive bids (i.e., never bid below $x - \epsilon$). The outsiders still have the problem of avoiding the winner's curse, only now they are competing against a bidder who knows the value of the item with certainty. In implementing this experimental design we conjectured that the existence of an insider might induce outsiders to bid more conservatively, thereby reducing or even eliminating the winner's curse. If the latter were correct, it would provide a possible insight into how adverse selection problems are avoided outside the laboratory, since many markets with a strong common value element (such as the used car market) have asymmetric information structures.

Table 7.9 reports data for high bidders from a series of asymmetric information auctions with inexperienced bidders, along with data from symmetric information auctions using the same subject population.[57] Data reported from the asymmetric information auctions are restricted to outsiders (the less informed) bidders. Profits for the outsiders are negative and significantly below zero for both values of ϵ, indicating that the winner's curse is alive and well in auctions with asymmetric information. Further, bids commonly exceed the expected value of the item conditional on winning the auction, indicating that these negative profits were not a result of unusually bad realizations of x_0.[58] Finally, it does not appear that the

Table 7.9. Asymmetric Information versus Symmetric Information in First-Price Auctions: Inexperienced Bidders

	$\varepsilon = 6$				$\varepsilon = 12$			
	Actual Profit[a]	Discount $(x-b)$[a]	Percentage of Auctions Won by High Signal Holder[a]	Frequency of $b_1 > E[x_o \mid X = x_1]$[b]	Actual Profit	Discount $(x-b)$[a]	Percentage of Auctions Won by High Signal Holders[b]	Frequency of $b_1 > E[x_o \mid X = x_1]$[b]
Asymmetric Information[c]	−3.67 (0.46)	−0.46 (0.53)	61.5 (8/13)	100.0 (12/12)	−2.71 (0.98)	−4.72 (0.89)	70.8 (17/24)	87.5 (21/24)
Symmetric Information	−3.38 (0.91)	−0.89 (0.780)	33.3 (5/15)	80.0 (12/15)	−0.61 (0.85)	−6.67 (0.65)	63.6 (28/44)	65.9 (29/44)
Difference: Asymmetric Information less Symmetric Information	−0.29 (1.07)	0.43 (0.99)	28.2	20.0	−2.10 (1.32)	1.95 (1.10)	7.20	21.6

Notes: $n = 7$. All data are for high bidders only.

[a] Mean value, with standard error of the mean in parentheses. Average profit in asymmetric information auctions is conditional on outsiders' winning the auction.

[b] Mean value, with number of auctions in parentheses.

[c] Data for less informed bidders only.

existence of insiders materially reduces the aggressiveness of outsiders' bids, as the average difference between outsiders' signal values and bids (the bid discount) is somewhat smaller, and the average losses somewhat larger, than in the corresponding symmetric information auctions.[59]

C. *The Winner's Curse in Other Market Settings*

The potential for a winner's curse is not limited to common value auctions. Anytime items of varying quality are sold, and there is asymmetric information between buyers and sellers, there is an adverse selection problem with a potential winner's curse. This section reviews experiments dealing with three different markets of this sort: bilateral bargaining games with asymmetric information, "blind bid" auctions, and two-sided auctions where product quality is endogenously determined. Although inexperienced subjects suffer from a winner's curse in all three settings, the strength of the winner's curse varies considerably, being quite persistent in the bilateral bargaining game and being eliminated after a few replications when product quality is endogenously determined. The latter outcome results from the strong incentive sellers have to dilute quality, which creates such a strong adverse selection problem that the market quickly unravels to the point that only "lemons" are sold, with buyers paying lemon's prices.

1. The Winner's Curse in Bilateral Bargaining Games with Asymmetric Information

Samuelson and Bazerman (1985) explored the winner's curse in a bilateral bargaining game with asymmetric information. This game adapted Akerlof's (1970) example of adverse selection in a market with lemons. Buyers (the acquiring firm) know that the target's value (v) is uniformly distributed in the interval [0, \$100], with the value to the buyer being $1.5v$. Buyers do not know v at the time they bid, but sellers do. Under this design, with sellers following the dominant strategy of only accepting offers that are greater than or equal to v, expected profit to the buyer for any *accepted* bid is negative, so the optimal bid is zero.[60]

The original Samuelson and Bazerman study consisted of a game where buyers bid one time only, and the experimenter acted as the seller. Comparing games with and without monetary incentives, the overwhelming number of subjects (92 percent and 93 percent, respectively) fell prey to the winner's curse, bidding positive amounts for the target firm.[61] The majority of subjects (59 percent with monetary incentives, 73 percent without) followed the naive strategy of bidding somewhere between the expected value of the target (\$50) and the expected value of the target in the hands of the acquiring firm (\$75), with most others bidding less. Samuelson and Bazerman attribute the modest downward shift in the bid distribution under monetary incentives to risk aversion.[62]

Ball, Bazerman, and Carroll (1991) extended Samuelson and Bazerman's design to allow for learning from repeated play. Each subject played a total of twenty trials, and players received feedback regarding the value of the company and how much money was made or lost following each trial. Players were given starting cash balances that were more than sufficient to cover any expected losses and were paid their net, end of experiment balance in cash. Figure 7.4 shows their results. There is virtually no downward adjustment in mean bids from the beginning to the end of the twenty trials, with bids over the first three trials averaging \$57 compared to \$55 over the last three trials. Only 7 percent of the subjects (five out of sixty-nine) learned to avoid the winner's curse during the experiment, defined as bidding either \$0 or \$1.00. Students with quantitative backgrounds (e.g., engineering degrees) did no better in avoiding the winner's curse than those with nonquantitative backgrounds (e.g., English degrees).

The Ball et al. experiment, like the Samuelson and Bazerman study, employed "realistic" instructions, subjects being told the following:

> In the following exercise you will represent Company A (the acquirer), which is currently considering acquiring Company T (the target) by means of a tender offer. (Ball et al. 1991, Appendix 1)

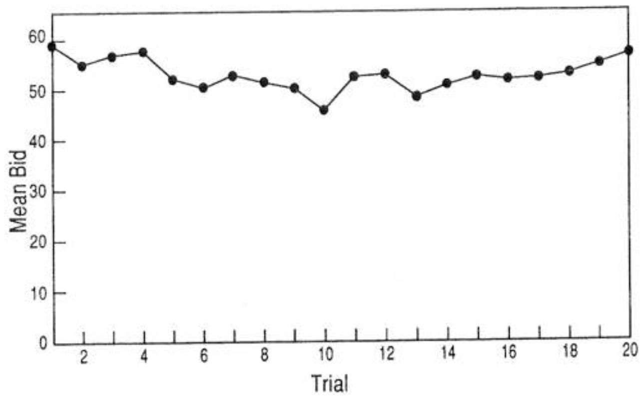

Figure 7.4. Average bids in bilateral bargaining game. *Source:* Ball, Bazerman, and Carroll 1991.

Further, subjects were explicitly told that the target company was expected to accept any offer greater than or equal to its value, but that

> acquiring a company is a neutral event—your performance will be judged only on the value of your assets at the end of this exercise. (1991, Appendix 1)

Nevertheless, for many experimentalists (myself included), embedding the game in a takeover context, particularly with MBA students, might be expected to generate overly aggressive bidding as a result of demand induced experimenter effects (see Rosenthal, 1976, chap. 11, for example). In response to this criticism, Ball (1991) has replicated the twenty-period design using value-free instructions, with no difference in the outcome: 5 percent of the subjects (two out of thirty-seven) avoided the winner's curse.

Cifuentes and Sunder (1991) have independently replicated the Ball et al. results using value-free instructions and forty trials in a session. In addition, under one treatment condition, v was distributed uniformly on the interval [\$10, \$100]. Since the value to the buyer was $1.5v$, this creates nonnegative expected profit (hence no winner's curse) for bids between \$10 and \$30, with an optimal bid of \$19 to \$20, generating an expected profit of \$0.28 per trial. Thus, buyers did not have to withdraw from active bidding to avoid the winner's curse. This is important since it might be argued that subjects bid positive amounts in the Ball et al. design out of boredom.[63] Nevertheless, Cifuentes and Sunder found few subjects avoiding the winner's curse: only two out of thirteen learned to bid 0 consistently with $v \in$ [\$0, \$100], and only 7 percent of all bids were between \$10 and \$30 for $v \in$ [\$10, \$100]. Overall, bids averaged well above the equilibrium prediction.

A Loser's Curse

Holt and Sherman (1994) conjecture that "Other possible influences, such as the thrill of winning, can lead to the same effect in bidding for objects of unknown value as the winner's curse." To test this conjecture, they modified the parameters

of the bilateral bargaining game to produce a "loser's curse," parameter values under which a naive bidder, who ignores the seller's incentive, bids *less* than is optimal and, consequently, wins less often than he should. This is done as follows: Let the value of the item to the owner be $v = 0.5 + V$ where V is uniformly distributed on the interval [0, 0.5]. As in Ball et al., the value to the acquiring firm is $1.5v$, and the owner only sells when the bid $b \geq v$. Under these parameters the optimal naive bid, the bid which maximizes expected value ignoring the seller's acceptance rule, is *less* than the optimal bid that conditions on the seller's acceptance rule, thereby generating a "loser's curse". In contrast let $v = 1.5 + V$ where V is uniformly distributed on the interval [0, 4.5]. This produces a standard winner's curse with the optimal naive bid greater than the optimal bid which conditions on the seller's acceptance rule.[64] Further, the expected loss from overbidding here exactly matches the expected loss from underbidding in the loser's curse.[65]

If the winner's curse is a result of the thrill of winning rather than a failure to account for the seller's acceptance rule, than underbidding will not materialize in the loser's curse, or if underbidding does occur, it will be substantially weaker than overbidding in the winner's curse. Results from this experiment show that average actual bids are remarkably close to the naive bid under both sets of parameter values. As such, Holt and Sherman's data indicate that the thrill of winning does not explain the winner's curse.

2. The Winner's Curse in "Blind Bid" Auctions

Consider the following extension of the bilateral bargaining game. Several buyers compete to purchase an item whose value (v) is common to all buyers and is uniformly distributed on some known interval. The seller receives no value from retaining the item and sells it in a first-price sealed bid auction. Further, the seller, who knows the exact value of the item, has an option, prior to the sale, of revealing the value of the item or concealing it (in which case the item is said to be "blind bid"). Buyers have no information other than what the seller provides them and the distribution of v.

If none of the sellers reveals the value of the item, buyers should bid the expected value, earning negative profits on low valued items and positive profits on high valued items. However, sellers have an incentive to reveal information on high valued items in order to receive full value for them. In this case buyers face an adverse selection problem for items that are blind bid. If buyers adjust to this adverse selection problem, the only sequential equilibrium is one in which the game completely unravels, with sellers revealing information on all but the lowest valued items and being indifferent between revealing or concealing information on the low valued items (Milgrom and Roberts 1986; Forsythe, Isaac, and Palfrey 1989). Failure to recognize this adverse selection problem results in a winner's curse, with buyers earning negative average profit on the blind bid items.

Forsythe, Isaac, and Palfrey (1989) report an experiment investigating blind bid auctions of this type.[66] Each experimental session consisted of a series of

auctions, so there was time for feedback and learning. Sellers were required to reveal truthfully the value of items that were not blind bid. Values of blind bid items were publicly announced at the end of each auction period.

Under these procedures, Forsythe et al. report clear evidence of a winner's curse in early auction periods. First, there was an adverse selection problem as seller's blind bid lower valued items right from the start.[67] Second, the winning bid was greater than the value of the item for 69 percent (59/85) of all blind bid items. This winner's curse was present in later auction periods as well, with buyers suffering losses on two thirds of the blind bid items. Further, sellers took advantage of the winner's curse as almost all low valued items (96 percent) were blind bid, although the theory predicts that sellers will be indifferent between blind bidding or revealing information on these low valued items.

More important, however, than the existence of a winner's curse in this experiment is the unraveling reported within each auction series as fewer and fewer items were blind bid over time: 44.2 percent of all items in the first ten periods compared to 28.1 percent from the eleventh period on. In addition, the value of blind bid items decreased monotonically over time, as the unraveling argument predicts, as did the prices paid for the items and the average losses on blind bid items.[68] These results are illustrated in Figure 7.5, which shows the value of the blind bid items and the amounts bid, by auction period, from two of Forsythe et al.'s markets.

These results lead Forsythe et al. (230) to conclude that

> the practice of blind bidding causes no difficulties once an equilibrium is obtained. After sufficient market experience, sellers only blind bid low quality items, and buyers, bids indicate that they had adjusted their beliefs properly.

In other words, most of the excess profit sellers earned on low valued items occurred in the initial auction periods prior to the unraveling (equilibrium) outcome. The relatively rapid reduction in bids and losses on blind bid items stand in marked contrast to the persistence of the winner's curse in the bilateral bargaining game. What seems partly, if not entirely, responsible for these differences is the role that sellers play. In the blind bid auctions sellers must sell the item to make any profit, and it is in their best interest to respond to buyers' naive bidding strategies by revealing information about higher valued items. This in turn creates an adverse selection problem on the blind bid items for which buyers do not fully compensate in their bidding (as they continue to suffer losses on two thirds of all blind bid items), but to which they partially adjust by lowering their bids. And this process repeats itself, as sellers now have incentives to reveal information on even lower quality items. Consequently, the market structure, with sellers having the opportunity and incentive to reveal information on lower valued items, speeds up the unraveling, which naturally limits the extent of the winner's curse in these markets. As will be shown in the next section, similar forces are at work in markets where quality is endogenously determined.

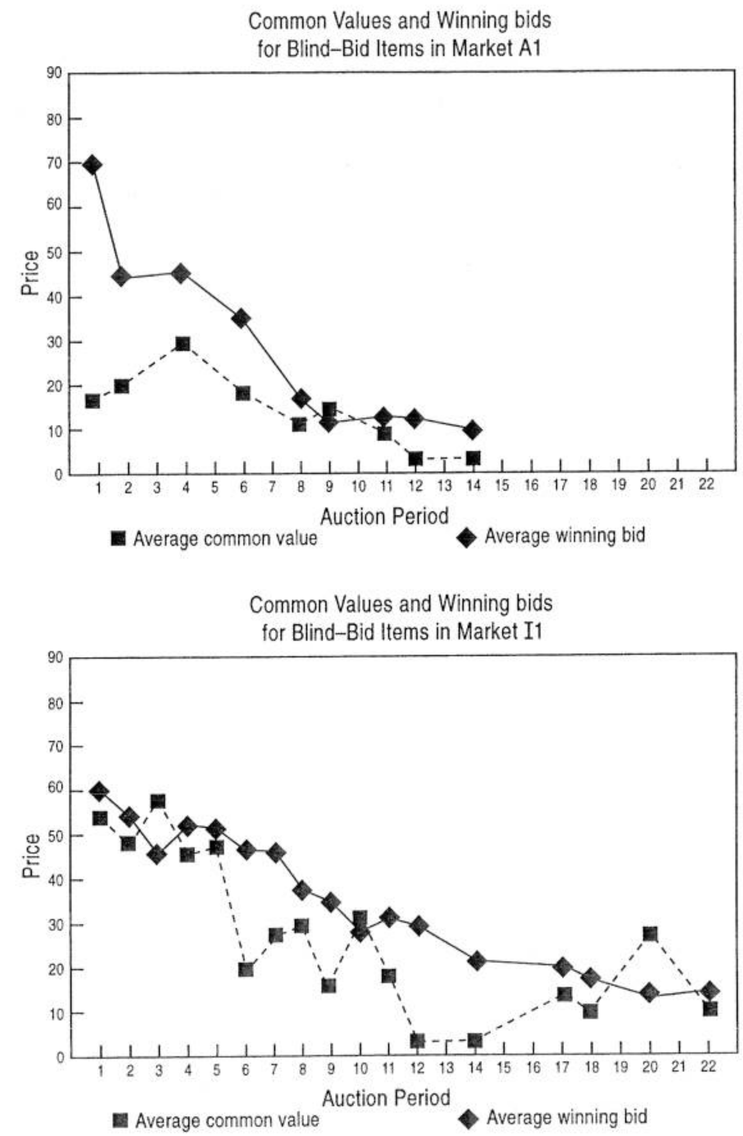

Figure 7.5. Common values and winning bids for blind bid items. Gaps in data points indicate that no items were blind bid that period. *Source:* Forsythe, Isaac, and Palfrey 1989.

3. Lemons and Ripoffs: The Winner's Curse in Markets with Quality Endogenously Determined

Lynch, Miller, Plott, and Porter (1986, 1991) study markets in which product quality is endogenously determined and cannot be observed by buyers prior to purchase (in all other experiments reviewed so far value has been exogenously determined). Sellers chose between producing a high or a low quality product where the marginal cost of producing the high quality product was greater than the low quality product. Buyers obtained greater value for the high quality product. It was common knowledge that it cost sellers more to produce "supers" (the high quality product) compared to "regulars." Markets were organized as continuous double auctions (see section III.C and Holt, chapter 5, section V) with several buyers and several sellers on either side of the market.

In markets where there were no enforceable warranties and sellers' identities were anonymous, so that they could not develop reputations for product quality, there is a potential "lemons" problem with sellers producing only low quality products. This in turn creates an adverse selection problem that, if not correctly anticipated, can result in large losses for buyers as a result of paying premium prices for what turn out to be low quality products (the winner's curse). As Lynch et al. (1986, 1991) demonstrate, inexperienced buyers fell prey to the winner's curse, even in this simplified setting, as sellers almost always produced a low quality product and buyers almost always paid prices greater than the highest possible redemption value for these low quality products, at least in the first market period. However, these large losses, in conjunction with sellers' continued delivery of low quality products, generated rapid adjustments on buyers' part so that prices were within 5 francs of those predicted by the lemons model by the fourth period in virtually all markets.[69]

Lynch et al. go on to explore the effects of warranties and sellers' reputations for correcting the lemons problem (warranties are effective, but reputation effects do not necessarily guarantee efficient market performance—see chapter 5, section IX). However, the relevant point to note here is that even in this simplified setting buyers initially fell prey to the winner's curse and that sellers' behavior, almost always supplying low quality products, appears to have been largely responsible for the rapid elimination of the problem.

4. Summing Up

In the bilateral bargaining game with asymmetric information buyers repeatedly suffer from a winner's curse as they fail to account for the adverse selection problem inherent in the seller's strategy of rejecting bids below the value of the item. It is clear at this point that the buyer behavior reported in the original Samuelson and Bazerman (1985) paper is robust and remarkably resistent to the feedback associated with repeated losses. (Factors that promote learning in this environment, and in common value auctions, are discussed in section II.D below.)

In both the blind bid auction experiment of Forsythe et al. and the lemons market of Lynch et al. there are potential adverse selection problems, which are realized. But the worst of the winner's curse effect is relatively short lived as the market unravels. In both cases the rather rapid elimination of the winner's curse is directly attributable to the market structure, in particular the actions sellers take in pursuit of their own self interest. This exacerbates the adverse selection problem and buyers' losses unless they make some adjustment in the right direction. In other words, in these markets the active role that sellers play seems largely responsible for the market unraveling and with it the elimination of the winner's curse. Further, new entrants, or infrequent participants, in markets such as this are likely to be protected from the strong winner's curse found in the start-up phase of the market by the fact that they unravel so quickly and are likely to stay unraveled with any sizable group of experienced bidders staying in the market.

D. Learning and Adjustment Processes in Markets with a Winner's Curse

This section is concerned with adjustments in bidding when sellers' actions do not contribute to eliminating the winner's curse—common value auction experiments and bilateral bargaining experiments with asymmetric information. Learning and adjustment over time are important phenomena in many experimental situations. Learning processes are of growing interest to theorists as well since it is clear that Nash equilibria are not established instantaneously, but rather are likely to evolve out of some sort of learning and adjustment process. In what follows, we report on what little is known about factors underlying subjects' adjustments to the winner's curse in these games.

Ball et al. (1991) and Cifuentes and Sunder (1991) report virtually no adjustment to the winner's curse in repeated bilateral bargaining trials (recall Figure 7.4). Indeed, the absence of any noticeable within session adjustments in both these studies is quite remarkable, at least to this reviewer. However, as Ball et al. show, behavior does adjust in response to (i) buyers playing the part of sellers and (ii) buyers returning for a second night of play.

To determine the effect of experience as sellers (the informed player) on performance as buyers (the uninformed player), Ball et al.'s subjects participated in two experimental sessions. The first session consisted of two parts: twenty trials in which subjects acted as buyers, followed by twenty trials in which they acted as sellers. In their role as sellers, subjects were told the value of the company they were trying to sell, given an offer and asked to accept or reject it (offers were taken from a representative subject's behavior as a buyer in an earlier experiment). Following each period, the computer calculated the sellers' earnings for that period and the period earnings for their imaginary buyer opponent. In the second experimental session, conducted several days following the first, subjects again played the role of buyers for twenty trials. As a control condition another group of subjects acted as buyers in an initial experimental session (of twenty trials), returning as buyers for a second experimental session several days later.

In the case where subjects acted as both buyers and sellers, the percentage of

learners (defined as subjects bidding zero from any particular trial until the end of the experimental session) jumped from 9 percent (four out of forty-four) to 37 percent (fifteen out of forty-one) between sessions 1 and 2 (a statistically significant increase).[70] Further, the mean bid for the non-learners changed from $51 to $34, suggesting that, although they had not learned to avoid bids with negative expected profit, they had learned to reduce their chances of losing money in any given trial and to reduce their losses conditional on the event of winning. It is interesting to note that these bidding adjustments took place at the beginning of the second experimental session, and that there was virtually no adjustment over time within the second session.

BBC report some adjustments in the control group as well. The percentage of learners increased from 6 percent (two of thirty-four) to 12 percent (four of thirty-four), although this difference is not significant at conventional levels. Further, the mean bid of nonlearners dropped from $50 to $34, almost the same as the reduction in bidding reported for the nonlearners who had played as sellers. Here, too, almost all the learning took place between the two experimental sessions, with minimal adjustments within the second session.

The fact that playing as sellers dramatically affected buyers' behavior squares nicely with one's intuition. The role reversal from buyer to seller literally forces buyers to act out the behavior of sellers, which should promote incorporating the informed party's (seller's) decisions into their bids after the experience. In other words, to the extent that the winner's curse results from buyer's failure to account for the adverse selection problem inherent in the seller's decision rule, playing the role of sellers promotes buyers accounting for this fact.

In common value auctions there are two distinctly different vehicles for bids adjusting over time. First, there is the possibility of "market learning" as bankruptcies drive out the more aggressive bidders and more aggressive bidders earn lower than average profits, self-selecting out of further experimental sessions. Second, there is the possibility for individual learning as subjects respond to repeated losses by bidding less. There is evidence that both factors are at work in common value auctions.

In Lind and Plott (1991), both the number of bidders and the size of ϵ were held constant throughout. Lind and Plott conclude that "the winner's curse persists with experience but the magnitude and frequency of losses decline with experience" (Conclusion 2). Clear evidence for this is found in Table 7.7, reported earlier. There, both loss frequencies and average profits are divided into ten period quartiles for each experimental session. As can be seen, with the notable exception of session 5, both the frequency and magnitude of losses decrease after the first ten trials.[71] To the extent that these adjustments are real rather than spurious, they must be the result of individual learning as no bankruptcies occurred in these sessions.[72]

Garvin and Kagel (1991) identify substantial adjustments in bidding in first-price auctions between inexperienced and experienced bidders (subjects who have participated in one previous first-price auction series). Table 7.10 reports some of their results. For example, with n = 4, the winning bid is *greater* than

Table 7.10. Effects of Experience on Bidding in First-Price Common Value Auctions

		Percentage of Auctions Won with x_1^a	High Bidders: Percentage of $b_1 > E[x_o \mid X=x_1]^a$	All Bidders: Percentage of $b_1 > E[x_o \mid X=x_1]^a$	Percentage of Auctions with Positive Profit^a	Average Actual Profit^b	Average Predicted RNNE Profit^b	Average Discount $(x-b)^{b,c}$
$n = 4$	Inexperienced	73.2 (30/41)	75.6 (31/41)	64.6 (106/164)	43.9 (18/41)	−1.32 (0.79)	5.01 (0.60)	4.60 (0.69)
	Experienced	66.3 (59/89)	34.8 (31/89)	29.5 (105/356)	58.4 (52/89)	1.37 (0.49)	4.32 (0.41)	8.07 (0.36)
$n = 6$ or 7	Inexperienced	56.7 (17/30)	76.7 (23/30)	58.2 (121/208)	30.0 (9/30)	−3.75 (0.89)	2.76 (0.53)	3.56 (0.97)
	Experienced	47.4 (9/19)	68.4 (13/19)	33.8 (45/133)	31.6 (6/19)	−0.32 (0.56)	2.93 (0.54)	5.53 (1.10)

Source: Garvin and Kagel 1991.
Note: ε = $12.
[a] Numbers in parentheses are raw data.
[b] Numbers in parentheses are the standard error of the mean.
[c] Data for $x \geq \underline{x} + \varepsilon$.

$E[x_o \mid X = x_1]$ in 75.6 percent of all auctions with inexperienced bidders compared to 34.8 percent with experienced bidders. Or, looking at the same behavior from a different point of view, for inexperienced subjects the average discount $(x - b)$ for winning bidders was $4.42, well below what is required to avoid the winner's curse (a discount of $7.20 with ϵ = $12), compared to an average discount of $7.44 for experienced bidders. This overly aggressive bidding translates into negative average profits for inexperienced bidders of $1.32 per auction period and losses in over 50 percent of all auction periods, compared to positive average profits of $1.37 and profits in 58.4 percent of all auctions for experienced bidders.

Garvin and Kagel (1991) identify two mechanisms behind the changes reported. First, a kind of market learning/self-selection effect is at work as bankrupt subjects were much less likely to return for the second auction series: 27 percent of bankrupt bidders (four out of fifteen) returned, while 86 percent of the solvent bidders (thirty-two out of thirty-seven) returned.[73] Further, subjects who failed to return for further experimental sessions (which includes bankrupt bidders), bid more aggressively, on average, as measured by differences in the average discount rate (the discount rate is defined as a bidder's signal value minus their bid, divided by ϵ).[74] Second, there were clear individual subject learning effects as inexperienced bidders bid less in response to actually losing money (experiential learning) and in response to potentially losing money had they applied their bidding strategy to the high bidder's signal (observational learning). Under some

conditions responses to these "hypothetical" losses were almost as large as the response to actually losing money. In some cases losing bidders who would have made money had their bid won the auction bid higher following auctions in which the high bidder made a positive profit. Although these "hypothetical" gain effects retard convergence to the RNNE (since bidding was above the equilibrium to begin with), they are smaller and less frequent in magnitude than the hypothetical losses, so the net effect of observational learning is to move bidding closer to the RNNE.

One noticeable difference between the common value auctions and the bilateral bargaining experiments is the within session adjustments to the winner's curse reported for inexperienced subjects: in the bilateral bargaining experiments there is virtually no adjustment, while reasonably large reductions in bids are found in the auctions. It is interesting to conjecture as to the basis for these differences. Ball et al. (1991) argue that there is a higher frequency of earning positive profits conditional on winning in the bargaining game (this happens 1/3 of the time with a naive bidding strategy) than in the auctions, so that the message, "bid less," comes through much more clearly in the auctions. Although this observation is true of the studies reported in Kagel et al. (1989), inexperienced bidders earned positive profits in 30 percent or more of the auction periods in Garvin and Kagel (Table 7.10) and in the first ten auction periods in Lind and Plott (Table 7.7). So this explanation appears suspect. A second factor to consider is that actual losses are often considerably larger in the auctions—they average around –\$1.30 in Table 7.10 and around –\$8.00 in the first ten periods in Lind and Plott—compared to –\$0.06 per period in Ball et al. and Cifuentes and Sunder (experimental dollars were converted into cash at the rate of 1 cent to the dollar in the bilateral bargaining studies). Perhaps subjects are more sensitive to these larger losses (recall section I.G).[75] Third, the observational learning reported by Garvin and Kagel in common value auctions is not available in the bilateral bargaining game where bidders only get to see the outcome of their own choices.[76] It remains to sort out carefully between these alternative explanations.

III. Additional Topics

Several additional topics are reviewed here. Section III.A reports the limited evidence on collusion in experimental auction markets. Section III.B relates the experimental findings to results reported from field studies. Sections III.C and III.D conclude with a brief discussion of two-sided (double) auctions and some other uses of auction market experiments reported in the literature.

A. *Collusion*

Most of the auction theory literature assumes that bidders act noncooperatively. This assumption may not be appropriate under some conditions, to the point that in field settings there is constant concern with the possibility of overt or tactic collusion (Cassady 1967, Graham and Marshall, 1987). With the notable exception of one auction series reported in Lind and Plott (1991), outright collusion among bidders has not been reported under standard experimental procedures. (The collusion reported in Lind and Plott was short lived and, apparently, unsuccessful, as there was always at least one defector.)

A number of significant obstacles must be overcome to mount a successful conspiracy in laboratory auction markets. Communication among bidders is typically not allowed, and subjects are usually brought together for a single auction session or, in the case of repeated sessions, group composition typically changes between sessions. As such, although there may be occasional signalling of a willingness to cooperate by an individual subject through, for example, submitting an unusually low bid, there is little chance for successful conspiracies to develop.

Isaac and Walker (1985) studied conspiracies through allowing bidders to discuss and coordinate bidding strategies. Discussions were allowed after bidders received their private resale values, but before they were required to bid. Isaac and Walker's announcement read as follows:

> Sometimes in previous experiments, participants have found it useful, when the opportunity arose, to communicate with one another. We are going to allow you this opportunity while we reset the computer between periods.
>
> There will be some restrictions.
>
> You are free to discuss any aspect of the experiment (or the market) that you wish, except that:
>
> (1) You may not discuss any aspect of the quantitative information on your screen.
> (2) You are not allowed to discuss side payments or physical threats.

A monitor was present to insure these restrictions were satisfied. A first-price (discriminatory) private value sealed bid auction was employed with four bidders and either one or three units of the commodity offered for sale, with feedback on auction outcomes consisting of either the winning bid (and that bidder's subject number) or all bids (with the associated subject numbers).

Isaac and Walker (1985) report active conspiracies developing in seven of twelve auction series, where a conspiracy is defined as all four bidders actively attempting to formulate and implement a collusive bidding scheme. Of the five remaining auction series, in two of them some bidders actively attempted to conspire, but at least one bidder openly refused to cooperate, and in the remainder the discussion never proceeded as far as an explicit conspiratorial scheme. Auction series with active conspiracies typically evolved to the point where the player(s) with the highest resale value(s) won the item with a minimum bid, with their co-conspirators bidding zero. Auction series with active conspiracies were stable, there being only three apparent cases of cheating. The results were ambiguous regarding Isaac and Walker's conjecture that conspiracies would be less successful if losing bids were not announced.

Dyer (personal communication) investigated tacit collusion in first-price private value auctions with three bidders. His treatment variable involved bidding in auctions with a fixed number of other bidders whose identity was known

compared to bidding in groups whose composition was randomly determined between auctions (the length of time these treatments were in effect was not announced, but they lasted a minimum of nine auction periods). The same set of subjects participated in three auction series, which took place over a two week period. Using a fixed effects regression model to analyze the data, his results are sensitive to both the regression specification and the data set employed. Fitting separate regressions to each auction series, he identifies a statistically significant effect in the expected direction (lower bidding with fixed identity) in the first auction series using a nonlinear time trend specification (some sort of time trend specification is necessary since bids as a proportion of resale values drift down over time—see section I.F above). But a linear time trend yields the opposite sign for this auction series (and is not statistically significant) and there are no significant treatment effects in the second and third auction series (the signs of the treatment variables are not consistent either). Fitting a single regression to the entire data set, a linear time trend specification shows a statistically significant effect in the expected direction, but the nonlinear time trend specification yields the opposite sign (which is not statistically significant). Adjustments in bidding reported across experimental sessions suggests that the nonlinear specification provides a more faithful representation of the data (see section I.F). Given these mixed results and the importance of the issue (is there a supergame effect when employing the same group of subjects within a given auction series) there is a clear need for more study.

My students and I have studied the effects of explicit opportunities for collusion in first-price common value auctions, implementing a variant of Isaac and Walker's procedures.[77] In common value auctions the experimenter faces the obvious problem of potentially large losses since x_0 varies over reasonably wide intervals (see section II.A). To solve this problem, a reserve price rule was established by drawing another signal at random from the interval $[x_0 - \epsilon, x_0 + \epsilon]$ and subtracting ϵ from it. A winning bid had to meet or beat the reserve price for the item to be sold. Two different information conditions were implemented, one in which both the reserve price rule and its realization was announced prior to bidding, and one in which only the price rule, but not its realization, was common knowledge. Two auction series were conducted, both beginning with five periods using the procedures described in section II.A, but with the addition of the reserve price rule. Subjects were then provided with an opportunity to communicate using an announcement similar to Isaac and Walker (1985), except that bidders *were* allowed to discuss and compare their private information signals. After ten auction periods, information conditions were changed from announcing the reserve price realization to no announcement in series 1, and from no announcement to announcing the reserve price realization in series 2.[78] These conditions were maintained for an additional ten auction periods.

Table 7.11 reports the results of these two auction series. In both cases, bidders suffered from the winner's curse under the no communication condition, which is not surprising as subjects had no prior experience with common value auctions. Collusion was active in both auction series when it was permitted, with a rotation

Table 7.11. Mean Profits in Common Value Auctions with Conspiracy

			Discussions			
	No Discussions		Reservation Price Announced		Reservation Price Not Announced	
Auction Series	Actual	Predicted[a]	Actual	Maximum Possible[b]	Actual	Predicted[c]
1	−2.54 (0.88)	1.66 (0.63)	4.85 (1.16)	4.86 (1.16)	1.21 (0.54)	2.05 (0.85)
2	−1.06 (0.50)	1.37 (0.52)	3.71 (1.06)	3.79 (1.06)	−0.07 (0.58)	0.00 (0.00)

Note: Numbers in parentheses are the standard error of the mean.
[a] Assuming RNNE bidding with no bid floor.
[b] Assuming bid floor provides focal point for conspiracy.
[c] Assuming optimal information pooling and risk neutrality in setting conspiracy price (see text).

rule (phases of the moon rule) used to determine the winning bidder (there was only one apparent break from the collusive arrangement, and that lasted for a single period). When the reserve price was announced it served as a focal point for the collusive price, with the winning bid being at or a few cents above it.

In cases where the reserve price was not announced, but where the rule was known, risk neutral bidders could maximize expected profit by pooling their information to establish an estimate of x_0, with the winner bidding $E[x_0] - \epsilon$. (Optimal pooling of information here requires averaging the low and high signal values to estimate $E[x_0]$.) In fact, the winning bid was well above $E[x_0] - \epsilon$, averaging \$1.56 (0.703) and \$1.95 (1.18) above it in the two auctions (standard errors of the mean are in parentheses).[79] From the information at hand it is not possible to tell if bidding above $E[x_0] - \epsilon$ resulted from a failure to properly pool information or from efforts to increase the chances of coming in above the unannounced reservation price (as might be expected if bidders are risk averse).[80] The net result was that actual profits were about half of what they would have been had the winner bid $E[x_0] - \epsilon$.

One interesting side product of this study is the realization that the reserve price, when announced, provides a focal point for collusive pricing should it exist. In contrast, announcing the price rule, but not its realization, results in risk neutral conspirators failing to purchase a number of items (half of them, on average, in our design) and possibly raising average prices on account of risk aversion. Although it is a robust prediction of auction theory that reserve prices will be used and announced, reserve prices, when they exist in practice, are typically not announced. Several elegant explanations of this apparent discrepancy between theory and practice have been offered in terms of noncooperative game theory

(McAfee and McMillan 1987c, Graham, Marshall, and Richard 1990). As seen here, an alternative explanation for this practice is the tendency for bidders to collude in auctions, a problem of considerable concern in practice (Cassady 1967), and the fact that when a reservation price is announced, it serves as a focal point for the collusive outcome.

B. *Comparing Results from Field Studies with Experiments*

There are important tradeoffs in studying auction theory using field as compared to experimental data. In an experiment, the researcher has full control of the auction structure so that, for example, a pure common value or private value auction can be constructed and the private information signals/valuations underlying bids are known. Failure to control these factors can make interpretation of field data problematic, as most actual auctions contain important private as well as common value elements, and the investigator must construct proxies for the private information signals/valuations underlying bids. Laboratory control also permits quite precise and demanding tests of the theory and facilitates identifying the causal factors underlying the behavior reported. On the other hand, field data have the advantage that they involve experienced professionals, with substantially larger amounts of money at stake than in the typical auction experiment, which may result in fundamentally different behavior. Wilson (1992, 261–62) summarizes the differences between field studies and laboratory experiments nicely:

> Empirical (field) studies must contend with less complete data, and few controls in the auction environment are possible. On the other hand, they have the advantage that the data pertain to practical situations in which the stakes are often large and the participants are skilled and experienced.

Despite these differences, results from field studies and experiments are complimentary in the sense that they both deal with the same phenomena, auction behavior. As such, both similarities and differences in behavior need to be identified (and, in the case of differences, hopefully reconciled) in order to enhance understanding. One purpose of this section is to make these comparisons.

A middle ground between field studies and experiments is to bring experienced professionals into the laboratory to participate in an experiment. Here, one presumably has the best of both worlds, strict control over the structure of the auction and experienced professionals to behave in it. Dyer et al. (1989b) did this in an experiment comparing the behavior of student subjects with construction industry executives in a common value offer auction. Behavior was found to be qualitatively similar as both subject populations fell prey to the winner's curse. This raises the puzzling issue of why presumably successful executives from the construction industry, an industry in which the competitive bidding process is often characterized as essentially a common value auction, could do so poorly in the laboratory.[81] The second purpose of this section is to provide summary results from a field study of the construction industry designed to answer this question.

1. Direct Comparisons between Laboratory and Field Data

In comparing empirical results from field studies with experiments, I rely heavily on the summaries of field work provided in McAfee and McMillan (1987a) and Wilson (1992). McAfee and McMillan note two unsurprising predictions of auction theory that have been confirmed using field data (primary sources underlying these conclusions are reported in parenthesis): (i) other things equal, a bidder with a higher valuation will submit a higher bid (in the case of offer auctions, a firm with lower costs submits a lower bid) (Gaver and Zimmerman 1977), and (ii) competition matters, so that the winning bid increases as the number of bidders increases (in the case of offer auctions, the winning bid falls) (Gaver and Zimmerman 1977; Brannman, Klein, and Weiss 1984). A third unsurprising prediction confirmed in field data is that in common value auctions better informed bidders make a higher rate of return than less informed bidders (Mead, Moseidjord, and Sorensen 1984; Hendricks and Porter 1988). All three of these predictions have also been confirmed in experimental data: (1) Bid functions estimated from private value auction experiments are strongly increasing in bidders' private valuations (Cox et al. 1988; section I.F, Table 7.4, above), while those for common value auctions are increasing in the buyers' signal values (Kagel and Levin 1986; Kagel et al. 1989). (2) Bid functions in first-price private value auctions are increasing in the number of bidders so that the winning bid must be increasing (section I.C), while bidders' profits are decreasing in common value auctions as n increases, so that the winning bid must be increasing as well (Kagel and Levin 1986). (3) In common value auctions with asymmetric information, less informed bidders' profits average around 25 percent of the informed bidders' profits (Kagel and Levin, unpublished data).

Considerable field work has been devoted to the study of outer continental shelf (OCS) oil lease auctions in the Gulf of Mexico for the period from 1954 to 1969. (OCS leases are often cited as the canonical form of the common value auction). Much of the research here, as in common value auction experiments, has focused on the existence and consequences of the winner's curse. Initial studies by petroleum geologists, who coined the term "winner's curse," claimed that winning bidders earned less than the market rate of return on their investments (Capen, Clapp, and Campbell 1971). Subsequent, more careful studies by economists yield more mixed results. Mead, Moseidjord, and Sorensen (1983) found after-tax rates of return to be somewhat less than average returns on equity for U.S. manufacturing corporations, concluding that

> they [the lessees] have historically received no risk premium and may have paid too much for the right to explore for and produce oil and gas on federal offshore lands. (43)

Note that different authors interpret the same results quite differently. Although Kagel and Levin (1986) argue that these results provide qualified support for Capen, Clapp, and Campbell's position, McAfee and McMillan (1987a) argue that they overturn these conclusions. Part of the difference in interpretation has to

do with whether investors require a risk premium for investing in oil and gas leases, with Mead et al. (1983) suggesting that they do. Others contend that large oil companies with access to capital markets and a diversified portfolio of leases would not be expected to earn risk premiums.

Hendricks, Porter, and Boudreau (1987) independently examined rates of return for OCS leases during the period from 1954 to 1969. Using somewhat different accounting procedures from Mead et al. (1983) they concluded that average realized profits were negative for auctions with more than six bidders, a conclusion remarkably similar to that reported in Kagel and Levin (1986). Hendricks et al. point out that these negative profits can be explained by nonoptimal bidding strategies that fail to account for the winner's curse, or equally, by adverse selection effects in estimating the number of bidders. That is, since most tracts receive less than six bids, and assuming that firms expect this, *ex post* profits will be less on tracts receiving more bids. Overall, Hendricks et al. conclude that "the data are consistent with both the assumptions and predictions of the [common value] model," allowing for bidders' uncertainty about the number of active bidders on each tract. As Wilson (1992) points out, however, this conclusion is stronger than in previous studies, where mixed results regarding profitability are often reported (Gilley, Karels, and Leone 1986).

Given the inconclusive nature of rate of return studies for OCS leases, Kagel and Levin (1986) use an anomalous finding reported in Mead et al. (1983, 1984) comparing rates of return on drainage with wildcat leases as suggesting important similarities between field and laboratory studies. A wildcat lease is one for which no positive drilling data are available, so that bidders have symmetric information. On a drainage lease hydrocarbons have been located on an adjacent tract so that there is asymmetric information, with companies who lease the adjacent tract(s) (neighbor(s)) having superior information to other companies (nonneighbors). The anomaly reported by Mead et al. is that *both* neighbors and nonneighbors earned a higher rate of return on drainage compared to wildcat leases. In other words, with asymmetric information, even the less informed bidders (nonneighbors) received a higher rate of return on drainage leases than on leases with symmetric information (wildcat tracts). Kagel and Levin (1986) rationalize this result by arguing that there is an important public information component to drainage tracts, and the public information may have corrected for a winner's curse that depressed rates of return on wildcat tracts. Details of this argument are outlined in chapter 1, section III.E and will not be repeated here.

A subsequent study of drainage lease auctions by Hendricks and Porter (1988), for this same time period, does not yield this anomalous result. Rather, Hendricks and Porter (1988) report the more conventional outcome that nonneighbors (those with inferior information) bidding on drainage tracts earned lower rates of return than firms bidding on wildcat tracts. However, updating their analysis to include production data from the 1980's, and more reliable production estimates prior to 1980, Hendricks and Porter (1992) obtain net rate of return estimates quite similar to Mead et al. (1984): *both* neighbors and nonneighbors earned a higher rate of return on drainage leases than the rate of return on wildcat leases.[82]

Anomalous results also emerge from field studies involving U.S. timber lease sales which have been conducted using both English and first-price auctions. Some have used these auctions to test the revenue-equivalence theorem: Do the two methods yield the same price on average? Using ordinary-least-squares regressions, Mead (1967) reports that sealed bid auctions had significantly higher prices than English auctions. Further study by Hansen (1985, 1986), who noted a selection bias caused by the way the Forest Service chose which auction to use, and corrected for it using a simultaneous equations model, found that although sealed bid auctions had slightly higher prices than the English auctions, the difference was not statistically significant, so that revenue equivalence could not be rejected. The puzzling part about these results is that there are strong common value (or at least correlated private value) elements to timber lease sales which should, in theory, result in English auctions raising more revenue. McAfee and McMillan (1987a) note this puzzle and go on to add

> The puzzle could be resolved by appealing to risk aversion of the bidders, but this remains an open empirical question. Given the sensitivity of the Revenue-Equivalence Theorem to its underlying assumptions, the theorem cannot be meaningfully tested until some way is found to test for independent private values against affiliated values. (727)

Experimental studies offer a possible alternative resolution to this anomaly. Common value auction experiments comparing sealed bid with English auctions show that English auctions do not consistently generate higher prices. In auctions where bidders commonly suffer from a winner's curse, as occurs with inexperienced bidders, sealed bid auctions consistently yield higher prices than English auctions (recall Table 7.8 from section II.B). Further, in auctions where bidders consistently earn positive profits, but still exhibit relatively strong traces of the winner's curse, as with moderately experienced bidders, the two auctions yield roughly the same prices (Kagel and Levin 1992). That is, in auctions where bidders suffer from traces of the winner's curse, the public information inherent in low bidders' drop-out prices tends to reduce rather than raise revenue (just as experimenter announced public information produces the same effect when bidders suffer from the winner's curse—see sections II.A and II.B above). Of course, to use this mechanism to resolve the anomaly in the timber sales field data requires postulating that bidders suffer from elements of the winner's curse.

It is worth adding here that in affiliated private value auction experiments, public information did not raise nearly as much revenue as predicted under risk neutrality (Kagel et al. 1987, reviewed in section I.D). Although risk aversion may partially explain these results, the data also show a sizable frequency of individual subject bidding errors associated with the release of public information, and these bidding errors inhibit the revenue raising mechanism underlying the theory. In short, the experimental evidence indicates that the revenue raising possibilities associated with affiliated value auctions are not nearly as robust as symmetric auction theory would lead one to expect, and the timber sale results are consistent with this fact.

2. Differences in Structure between Laboratory and Field Auctions

Dyer and Kagel (1992) address the question of why experienced construction industry executives fell prey to the winner's curse in laboratory offer auctions, as reported in Dyer et al. (1989b). They focus on two possibilities, which are not necessarily mutually exclusive. One is that the executives had learned a set of situation specific rules of thumb which permit them to avoid the winner's curse in the field, but which could not be applied in the laboratory. The second is that the bidding environment created in the experiment, which is based on theoretical work, is not representative of the environment encountered in the field.

Evidence supporting this first possibility emerges from interviews with contractors. In these interviews it is clear that an important determinant of the risk associated with bidding a job involves the architect/owner. The architect/owner's reputation for admitting mistakes, or ambiguities, in construction plans, and their willingness to implement functionally equivalent construction outcomes using alternative, and cheaper, construction techniques than originally specified, play an important role in the cost estimates assigned to specific job components, as well as the markup assigned to the total cost estimate.[83] In addition, firms tend to specialize in different types of construction projects (or at least their estimators do). Experienced contractors pride themselves on their familiarity with building different types of structures and figure their estimates to lie within a rather narrow band of the true value. This familiarity is based on past experience, to the point that in one bidding session I sat in on, the firm had just completed a similar building designed by the same architect. At one point, when in doubt on the cost estimate to assign to a particular component of the job they were bidding on, the bid team simply pulled up records from the recently completed job and filled in the missing numbers.[84] Needless to say, the contractors did not have these situation specific rules to rely on when bidding in the laboratory.

Evidence that the field environment differs in important ways from theoretical specifications operationalized in the laboratory is summarized in Table 7.12, which shows the distribution of bids on a particular job, measured in terms of deviations from the low winning bid. The first thing the reader will note is that the low bid is some $30,000 *below the low winning bid*. This was the result of a bidding "error" which resulted in the original low bidder withdrawing his bid without penalty after the bids were announced.[85] Standard auction theory does not account for such possibilities.[86] The second thing to note is the small difference between the winning low bid and the second lowest bid ("money left on the table"), less than 1 percent of the low bid. This difference is minuscule and indicative of the relatively small differences between the low bid and the second lowest bid characteristic of much of the industry, which averaged around 5 percent for the sample of jobs analyzed in Dyer and Kagel. By way of contrast, Hendricks et al. (1987) report money left on the table from OCS leases averages around 50 percent of the winning bid.[87] This implies that there is much smaller scope for the winner's curse to express itself in the branch of the construction

Table 7.12. Bids by Firm Measured in Terms of Deviation from Low Winning Bid

Firm	Deviaton from Low Bid (dollars)	Deviation as a Percentage of Low Bid	Firm	Deviaton from Low Bid (dollars)	Deviation as a Percentage of Low Bid
1	–30,000[a]	–0.71	8	105,000	2.47
2	0	0.00	9	142,000	3.33
3	32,000	0.75	10	144,000	3.38
4	64,000	1.50	11	155,000	3.64
5	74,600	1.75	12	166,000	3.90
6	87,679	2.06	13	183,000	4.30
7	90,000	2.12	14	564,000	13.25

[a] Mistake in bid and let out of bid. Second-highest bidder got the job. Mean bid $4.38 million.

industry in which these executives worked. As such, private value elements, such as the firm's overhead and the amount of idle resources anticipated, often play an important role in determining the low bidder.[88]

3. Summing Up

Unsurprising predictions of auction theory regarding higher winning bids with more bidders, and higher profits for bidders with superior information in common value auctions with asymmetric information, are confirmed in both field and laboratory studies. There were reports of a winner's curse in early OCS oil lease auctions, but these findings have been subject to considerable dispute, as have other citings of the phenomena. Anomalies identified in laboratory experiments have parallels in field data, but alternative explanations for the field data are available as well. Professional bidders from the construction industry fell prey to the winner's curse in a laboratory offer auction. Reasons suggested for this are that (i) learning tends to be situation specific, and the experiment stripped away many of the contextual clues the professionals employ in field settings and (ii) the construction industry has private value and repeated play elements that were not present in the experiment, which mitigate the winner's curse.

C. *Two-Sided Auctions*

Experimental investigations of two-sided (double) auctions have commonly employed stationary aggregate supply and demand schedules (schedules that remain constant across auction periods). This is in contrast to one-sided auction experiments where buyers' valuations/information signals are randomly drawn each auction period. Two-sided auctions, particularly continuous double auctions, have

been used in this way to investigate questions in industrial organization theory (see chapter 5) and in price formation in asset markets (see chapter 6). They have typically not been used to investigate theories of double auction markets in the same way that one-sided auction experiments have. The primary culprit here is that, until recently, there have been no clearly articulated theories of double auction markets in terms of Bayesian Nash equilibria of games with incomplete information (see Wilson [1992] for a review of recent research along these lines), in large measure because of the difficulties in modeling strategic behavior on both sides of the market.[89] Two preliminary experimental studies have, however, investigated recently developed Bayesian Nash equilibrium models of double auction markets. Both involve private value auctions with buyers' valuations and sellers' costs drawn randomly each auction period from known distributions and with buyers and sellers limited to trading a single unit.

Cason and Friedman (1993) study continuous double auction (CDA) markets. A continuous double auction (CDA) market permits trade at any time in a trading period, with buyers and sellers free to continuously update their bids and offers. Cason and Friedman test three CDA models concerned with price formation *within* trading periods. In order of decreasing rationality, these are Wilson's (1987) waiting game/Dutch auction model (a sequential equilibrium model in which transactions occur between the highest valued buyer and lowest cost seller remaining in the market), Friedman's (1991) Bayesian game against nature (where traders ignore the impact of their own current bids and asks on subsequent offers by others but employ Bayes' law in updating bids and offers), and Gode and Sunder's (1993) budget constrained zero intelligence (ZI) trading model (buyers bid randomly between their valuation and zero, and sellers offer randomly between their cost and the upper bound of the cost distribution).[90] The contrasting predictions looked at concern serial correlation between successive transaction prices, the nature of successive bid and ask sequences culminating in a transaction, the extent to which early transactions occur between high value buyers and low cost sellers, and the extent to which gains from trade are exhausted (efficiency).

Two data sets are used to test these predictions, only one of which strictly satisfies the theoretical requirements of all the models (buyers and sellers trade single units with new random valuations in each trading period). This one data set consists of three experimental sessions with four buyers and four sellers and inexperienced traders. Tests for serial correlation in transaction prices can not reject the null hypothesis of zero correlation, consistent with Wilson's (1987) waiting game model. However, the few traders involved result in only two or three transactions within each market period, so that the power of this test is quite low. Changes in successive bids and offers favors Friedman's Bayesian game against nature as successive improvements in bids (offers) culminating in a trade are most often made by different buyers (sellers) rather than the same buyer (seller) as the waiting game model predicts (ZI traders exhibit no consistent pattern with respect to this characteristic). Higher valued buyers and lower cost sellers tend to trade sooner, but the rank order correlations between buyer (seller) valuation and transaction order are weak and fall within the range predicted by the ZI algorithm. Efficiency levels are uniformly high, averaging 93 percent, which is a little lower than found in ZI simulations. Cason and Friedman are keenly aware of the data limitations underlying their analysis, calling for more experiments, particularly those with experienced traders, so that the common knowledge assumptions underlying Wilson's waiting game model are more likely to be satisfied.

Kagel and Vogt (1993) provide some initial data for Satterthwaite and Williams (1989a, 1989b, 1993) buyers' bid double auction (BBDA) markets.[91] The BBDA is a clearing house trading mechanism in which buyers and sellers submit a single bid and offer each trading period on the basis of which the market clears. The BBDA is a particularly attractive model to investigate since it makes a number of strong predictions: sellers have a (nontransparent) dominant strategy of offering at cost, given a uniform distributions for costs and valuations symmetric risk neutral buyers bid a simple proportion of their valuation which is increasing in the number of traders, and Satterthwaite and Williams have computed the exact increase in the expected gains from trade as the number of traders increase, assuming a symmetric RNNE. The latter provides the focus for Kagel and Vogt's experiment, as efficiency is predicted to increase from 93 percent in auctions with two buyers and two sellers, to near 100 percent in auctions with as few as eight buyers and eight sellers.

Two data sets were used to test these predictions, only one of which strictly satisfies the structural requirements of the theory. This one data set consists of two auction sessions with inexperienced bidders and one with experienced bidders. Against the benchmark of the symmetric RNNE model's predictions there were a number of shortcomings. Sellers did not adhere to the dominant strategy of offering at cost, which was not unexpected given results reported for the simpler one-sided second-price auctions and multiple unit uniform price auctions (see section I.B.2). Buyers bid well above the symmetric RNNE, much like the overbidding reported in one-sided first-price auctions (section I.G). Efficiency failed to increase significantly as the number of traders increased; rather, it increased modestly on average and the direction of change was erratic across experimental sessions. The fact that efficiency failed to increase can be largely attributed to buyers bidding substantially more than predicted in auctions with few traders and to there being practical upper bounds on how much bids can increase without exceeding buyers' valuations. However, efficiency measures were uniformly high, averaging 88 percent in markets with two buyers and two sellers and 90 percent in markets with eight buyers and eight sellers. This is substantially higher than predicted for ZI traders under the BBDA and consistently higher than predicted under a fixed price rule.[92] Kagel and Vogt conclude by noting the need for further research exploring the effects of instructions designed to better explain the somewhat complicated trading rules underlying the BBDA and the effects of computerized sellers, which may make for a more stable environment for learning to take place. There is certain to be considerably more research investigating the properties of two-sided auctions in terms of games with incomplete information in the future.

D. *Other Auction Studies*

Experimentalists have used auctions to investigate a variety of topics, not all of which could be covered in detail here: Cech et al. (1987) and Guler and Plott (1988) have used auctions to study incentive contract issues of the sort involved in military procurement contracts, Palfrey (1985) has used auctions to investigate the efficiency and distributional consequences of the common practice of selling a variety of different items in "bundles," and Pitchik and Schotter (1988) have investigated a budget-constrained, sequential auction with the primary purpose of testing the predictive power of trembling hand perfect equilibria. Auction markets serve as an ideal vehicle for experimental studies of these and a variety of other issues. As Wilson (1992) notes, auctions are apt subjects for applications of game theory because they present explicit trading rules that largely fix the "rules of the game." These well-specified rules of the game serve to fix the rules and design of an experiment. In addition, the strategic behavior that can be modeled using auctions is, more often than not, of practical importance as well. This promises to spur a number of new experimental applications in the future.

IV. Conclusions

Studies of experimental auction markets have been going on for over ten years now, paralleling the profession's interest in the theoretical properties of these markets. This research has established several facts about behavior relative to the theory.

In private value auctions the revenue-equivalence theorem fails. Bids in first-price auctions are higher than in Dutch auctions, and bids in second-price auctions are higher than in English auctions. This failure of the revenue-equivalence theorem may result from response mode effects with the sealed bid auctions focusing attention on the price dimension, and the open auctions focusing attention on earnings, thereby generating small but consistent differences in bids. This failure of the revenue-equivalence theorem appears to represent a simple bid level effect rather than a more fundamental breakdown in Nash equilibrium bidding theory as the theory correctly predicts the directional relationship between bids and valuations, and the directional effects of changing numbers of bidders, under different bid price rules. In addition, the Nash equilibrium bidding model performs better than sophisticated ad hoc bidding models in first-price auctions with affiliated private values. Still unresolved is the debate regarding the relative importance of risk aversion versus bidding errors in organizing deviations from the RNNE model's point predictions in first-price auctions.

Nash equilibrium bidding theory performs much worse in common value auctions. Common value auctions are substantially more complicated than private value auctions as they incorporate a difficult item estimation problem in addition to the strategic problems involved in competitive bidding. In common value auctions inexperienced bidders suffer from a winner's curse, falling prey to the

adverse selection problem inherent in these auctions. Overbidding here does not just involve a bid level effect but represents a more fundamental breakdown in the theory resulting in reversal of a number of important comparative static predictions: bidding does not decrease in response to increased numbers of bidders in second-price auctions as the theory predicts, and public information about the value of the item reduces, rather than raises, revenue in the presence of the winner's curse. This perverse effect of public information in the presence of a winner's curse extends to English common value auctions as well. Experienced bidders eventually adjust to the adverse selection forces inherent in common value auctions, but this largely reflects situation specific learning and market selection effects rather than theory absorption, so the ability of the theory to correctly predict bidding in field settings remains an open question.

The winner's curse extends to other market settings as well, being particularly robust in bilateral bargaining games with asymmetric information. However, blind bid auctions and markets where quality is endogenously determined unravel rather quickly, largely eliminating the winner's curse. What these markets add that the auctions and bilateral bargaining games lack is the dynamic interaction between sellers pursuing their own self interest, which exacerbates the adverse selection problem, and bidders adjusting in the right direction to the adverse selection effect. This process repeats itself until the market unravels, resulting in only lemons being sold or only low valued items being blind bid and with buyers largely adjusting to these extreme adverse selection effects.

Notes

Research support from the Economics and Information, Science and Technology Divisions of NSF, the Sloan Foundation and the Russell Sage Foundation are gratefully acknowledged. Special thanks to my coauthors Dan Levin and Ron Harstad for the many enlightening discussions on auction theory. Helpful comments on earlier drafts of this paper were received from Sheryl Ball, Raymond Battalio, Mark Isaac, John Ledyard, Dan Levin, Jack Ochs, Charles Plott, Rob Porter, Alvin Roth, Vernon Smith, James Walker, Elmar Wolfstetter, and participants at the Conference on the *Handbook of Experimental Economics*. Susan Garvin provided valuable research assistance throughout. I alone am responsible for errors and omissions.

1. There has been some experimental interest in one-sided auctions where bidders demand more than one unit of the good (Smith 1967, Belovicz 1979, Miller and Plott 1985; Grether, Isaac, and Plott 1989). There are limited theoretical results for such multiple unit auctions (see Forsythe and Isaac [1982] and Maskin and Riley [1989] for theory developments) and these have yet to be subject to systematic experimental investigation.
2. In the case of tied high bids, the winner is selected using a random tie breaking device.
3. These papers represent the first experimental studies of the IPV model with a clear theoretical focus. Frahm and Schrader (1970) conducted an earlier experimental study but employed a multiple unit sequential auction in which bidders did not know the number of units to be auctioned off (the latter was a stochastic variable whose mean and distribution was unknown). Frahm and Schrader offer no theoretical solution for their design.
4. Assuming some bidder heterogeneity, or random errors in bidding, efficiency measures will be sensitive to the support from which resale values are drawn and the number of bidders: other things equal, the larger the difference in $(\bar{x} - \underline{x})$ and/or the smaller n is, the higher the

average efficiency levels reported. Comparative efficiency measures are most meaningful when measured across auctions holding these variables constant.

5. Expected profits conditional on winning the auction were \$1.20 (Cox et al. 1982), \$1.50 (Cox et al. 1988) and \$7.50 (Dyer et al. 1989a).
6. I am grateful to my colleague Dan Levin for pointing this out to me. Cox et al. (1983a, 217) acknowledge this point as well.
7. Kagel and Levin (1993) report minimal learning on bidders' part as well, even following the occasional out-of-pocket losses that result from overbidding.
8. The Van Huyck, Battalio, and Beil results are part of a training exercise prior to participation in a more elaborate two stage game. The results in question have not been published, but they find some initial bidding above the dominant strategy price, similar to Kagel et al. (1987), with rapid convergence to the dominant strategy price (personal communication).
9. Several auction series with experienced subjects screened bidders to exclude those with the greatest absolute deviations from the dominant bidding strategy in the uniform price auctions, so the data from these auctions, must be treated with some caution.
10. A similar hysteresis effect is reported in Cox et al. (1985a) for multiple unit auctions. Bidding in an initial sequence of discriminatory auctions produced significantly lower bidding in subsequent uniform price auctions compared to bidding in an initial sequence of uniform price auctions. However, Harstad (1990) finds that prior experience need not necessarily affect bidding. In auction series involving several first-price auctions, followed by several second-price auctions, followed by several first-price auctions, bidding in the second series of first-price auctions is indistinguishable from the first series.
11. The expected cost of typical deviations from the RNNE found in first-price auctions is discussed in Harrison (1989) and section I.G.2 below.
12. My views on this issue have been heavily influenced by my colleague Dan Levin.
13. According to this literature the weight attached to a dimension increases when the dimension is psychologically more "compatible" with the response mode. Unfortunately, compatibility is not usually defined in advance, but is considered datum to be extracted from the problem. This does not obviate the apparent robustness of the effect. Tversky, Slovic, and Kahneman (1990) provide strong evidence for response mode effects underlying preference reversals. However, the exact cause of the phenomena is still a hotly debated topic (see the Introduction, section III.F.1, and Camerer, chapter 8, section III.I).
14. The crucial question subjects must pose to recognize the dominant strategy is do I want to win *as a result of* increasing my bid? The claim is that this question does not immediately leap to mind when bidding in a second-price auction.
15. Under the dual market technique each subject places two bids based on the same resale value (x), one in a market of size 5 and one in a market of size 10, with the market from which payment is made determined randomly after both bids are made. The dual market technique is as close as one can come to operationalizing the *ceteris paribus* conditions underlying the comparative static implications of economic theory. It has the advantage of controlling for extraneous variability in outcomes resulting from variations in private valuations across market conditions, as well as controlling for between subject variability, as the same bidder has the same valuation for the item in both markets. Assuming expected utility maximization on bidders' part, and randomly determining which market to pay off in, the optimal strategy in each market is independent of the other markets. The first extensive applications of this technique are reported in Smith (1980) and Palfrey (1985), but they did not investigate its properties relative to alternative procedures. Battalio et al. (1990) report extensive tests of whether the technique distorted bids in markets with changing numbers of bidders. They found no effects. Kagel et al. (1987) report less extensive tests of the technique in auctions with affiliated private values. They report no distortive effects as well.
16. In these models the number of bidders is determined exogenously and is not a function of resale values or the underlying support of the private values. This sort of situation would hold if, for example, auction participation were by invitation only, with the seller choosing to reveal or conceal the number of invitees.
17. See Milgrom and Weber (1982) for a formal definition and properties of affiliation.

18. Again, the intuition here is straightforward: Consider an extremely risk averse bidder who, under private information conditions, bids very close to his valuation out of fear of losing the object to rivals. Announcement of x_0 will, on average, assuage his fears, resulting in lower bidding and lower average prices.
19. Allowing for bidders learning about the experiment by throwing out the first three auction periods, high bids still averaged \$0.10 *below* the RNNE with $\epsilon = \$6$, but the null hypothesis of risk neutrality can no longer be rejected at conventional significance levels.
20. Bidders also show increasing relative risk aversion in the corresponding IPV auctions with x_0 announced.
21. Examination of bid tabulation data for construction contractors shows considerable diversity in the zeal with which different contractors obtain these data (Dyer and Kagel 1992). Records from at least one firm in our sample routinely show fairly complete information on bids from private sector jobs, with this information obtained from the architect and/or rival bidders. Many firms simply report "not known" for the distribution of bids on such projects. Section I.A characterizes the diversity in laboratory practices regarding price information feedback.
22. Stigler (1964), in a well known article, discusses the role this information can play in policing bid rings. For more on this see section 3.1 below.
23. On a more formal level, Cox et al. (1988) examined individual subject regressions in which they included a linear time trend variable. For 20% of the subjects the time trend coefficient is significant, with the significant subjects evenly split between increasing and decreasing their bids over time, suggesting no systematic within experiment learning effect.
24. The fixed effects regression model employs a different intercept coefficient for each subject to capture between subject variability in the data, but assumes identical coefficient estimates for the independent variables employed in the regression. Dyer also tested for tacit collusion in groups whose composition remained fixed between auction periods compared to groups whose composition was randomly determined in each auction period. These results are reported in section III.A.
25. The regression results reported here are new. I thank the authors of the original studies for providing me with access to their data.
26. The exchange in the *American Economic Review* fails to capture the intensity of the passions raised by this issue as the editors toned down the language on all sides for the published commentaries.
27. Cox et al. (1982, 1988) were unable to obtain an analytic solution for the bid function in cases where x would produce bids above b^*.
28. Cox et al. (1988) offer an extended version of CRRAM to account for the large number of bid functions with significant intercept values. The model incorporates a utility of winning the auction and a minimal income level (threshold) bidders aim for in winning the auction. When the utility of winning exceeds the income threshold, the bid function has a positive intercept; when the income threshold exceeds the utility of winning, the bid function has a negative intercept. To test this model Cox, Smith, and Walker conducted auctions with lump sum payments and lump sum charges for winning the auction. The prediction is that lump sum payments should produce higher intercepts than no payments and no payments should yield higher intercepts than lump sum charges. The intercepts line up as predicted for 78% of the pairwise comparisons reported. These experimental manipulations are of inherent interest, irrespective of any issues regarding CRRAM, since lump sum payments or charges should produce an intercept shift under RNNE bidding, as well as a variety of other model specifications, regardless of the existence of an income threshold for winning the auction or a utility of winning.
29. This result was obtained both in auctions with affiliated private values as well as in auctions with x_0 announced.
30. Kagel and Levin's (1985) tests were confined to IPV auctions (auctions with x_0 announced, see section I.D). This rejection rate is well above the 22% frequency with which Cox et al. (1988) report statistically significant intercept values in their regressions, but only slightly above the 38% frequency of nonlinear bid functions reported. Cox et al. (1992, 1407–8)

object to these results citing a working paper by Cox and Oaxaca (1992) in which they reanalyzed Kagel and Levin's data, introducing a time trend to account for possible changes in bidding over time. The Cox and Oaxaca analysis shows that 54% of the bidders violate the restrictions implied by CRRAM (which is within the range of rejections reported in Kagel and Levin under various assumptions regarding b^*). Cox and Oaxaca also develop and apply a spline-function technique that uses all the data in Cox et al. (1988) for $n = 4$ to estimate both parts of the CRRAM bid function. Nearly 50% (nineteen out of forty) of the individual subject regressions violate the CRRAM restrictions (at the 10% significance level or better; Cox and Oaxaca table 2). Nevertheless, Cox et al. (1992, 1401) claim that "the results [from Cox and Oaxaca] indicate that CRRAM organizes all of the data quite well."

31. This can be seen as follows. Normalizing $\underline{x} = 0$, profit for the high bidder under the RNNE is $x_i - [(n-1)/n]x_i$, while under CRRAM profits for the high bidder are $x_i - [(n-1)/(n-r_i)]x_i$. Assuming that bidders with the high resale value win the auction, as happens in over 80% of the auction periods, profits for CRRAM bidders as a percentage of the RNNE are $[1-(n-1)/(n-r_i)]/(1-(n-1)/n]$, which for positive r_i is increasing for decreases in n. Cox et al. (1992, 406) object to these tests on the grounds that (i) the "analysis is based entirely on the use of a linear bid function" but the data include observations from the upper part of the bid function (the non-linear portion according to CRRAM) and (ii) that the simple formula used here "has no valid implication for this question" since "correct expected profit calculations for these bid functions involve the use of order statistics." Objection (i) has not stopped Cox, Smith, and Walker from using *only* winning bids in testing CRRAM (for example Cox et al. 1983b). In fact, one of the motivations for the tests reported in Kagel and Levin (1985) was the weakness of the test reported in Cox et al. (1983b) which used only winning bids. Cox et al. (1992) have provided no counter example or details underlying objection (ii), which seems misdirected in this case.

32. In spite of misgivings about CRRA, my colleagues and I have had occasion to employ it when working with risk averse Nash bidding models. The reason for this is quite simple: CRRA provides one of the few closed form solutions available for characterizing risk averse Nash bidding in first-price auctions; hence it can be useful in providing some idea of the impact of risk aversion on auction outcomes. See Kagel et al. (1987) for just such an application.

33. Cox et al. (1992, footnote 14) argue that these negative correlations are to be expected in the context of an effort cost model of decision making: "This is because, in the model, the higher expected utility when values are higher results in increased opportunity cost of a nonoptimal decision; thus more decision effort is expended." This observation correctly captures the spirit of our argument and the main message, as I understand it, of Harrison's (1989) original argument.

34. Cox et al. (1992) object to these calculations on the grounds that since part of the data come from experiments using dual markets forgone earnings should be determined by multiplying through by the probability of being paid in a market. Kagel and Roth (1992) did not do this, noting that similar results are found independent of the number of dual markets employed. What the data appear to be telling us is that payoffs matter, but that there may be a threshold or context effect. In other words, bidders with high valuations put more effort and care into formulating bids in (all) dual markets (since they know there is a high probability of winning in one of the markets) than the effort and care expended with a low valuation in a single market with a comparable overall probability of winning.

35. Cox et al. (1992, 1401) argue, "Again, as in first-price auctions, this anomaly could be addressed with a utility of winning model (or a decision-cost model as in Smith and Walker)." It is hard to see how a utility for winning can organize this anomaly given that overbidding does *not* occur in English auctions (recall section I.B.2 above).

36. As Harrison (1992) points out, the payoff function under the Becker et al. procedure is quite flat, probably flatter than in Cox, Smith, and Walker's first-price auctions. Nevertheless, this is a nice example of the kind of internal consistency test that can be applied in the presence of inherently flat payoff functions.

37. For example, instead of paying $[x - b(x)]$ dollars to the winning bidder, the winner is paid $[x - b(x)]$ lottery tickets. The winner then participates in a post auction lottery in which she receives y_1 dollars with probability $[x - b(x)]/x$ and y_2 dollars ($y_1 > y_2$) with probability $1 - [x - b(x)]/x$. Losing bidders in the auction all receive y_2 dollars as well.

38. Harrison (1989) reports experimental treatments using a similar design with similar results.

39. Walker et al. note that they brought twelve of these inexperienced bidders back to play again without observing any further movement toward risk neutral bidding.

40. Using nonparametric test statistics applied to market prices, Rietz's first-price dollar auctions all resulted in price distributions lying significantly above the RNNE prediction, thereby demonstrating that in the absence of controls for risk aversion, subjects bid as if they were risk averse under his procedures.

41. Although this response may apply to multiple unit discriminatory and third-price auctions, it is of questionable relevance to second-price auctions.

42. Prasnikar and Roth (1992) (see Roth, chapter 4 in this text) report a common value auction in which there is no uncertainty about the value of the item. Prices converge to the value of the item within a few auction periods.

43. See Kagel and Levin (1986) for signals outside the interval $\underline{x} + \epsilon \leq x \leq \bar{x} - \epsilon$.

44. $E[x_0 \mid X = x_1] = x - \epsilon(n-1)/(n+1)$ for $\underline{x} + \epsilon \leq x \leq \bar{x} - \epsilon$

45. Roth (1988; chapter 1, section III.E) reviews earlier experimental work on the winner's curse where bidders were required to come up with their own estimates of the value of the item.

46. In the naive bidding model subjects act as if they are in an auction with affiliated private values, as they ignore the adverse selection problem but discount their bids relative to their signal values out of strategic considerations (following the dictates of equation (6) in section I.D). The naive bidding model insures negative expected profit whenever $n > 3$.

47. These calculations rely on the fact that $Y = [2\epsilon/(n+1)] \exp[-(n/2\epsilon)(x - (\underline{x} + \epsilon)]$ in equation (9) and assumes $\underline{x}$ is \$25. For RNNE bidding, the maximum possible losses occur when $x = \underline{x} + \epsilon$ and are equal to Y (Kagel and Levin 1991).

48. Literally, this is only a local result. However, in the simulations reported in Kagel and Levin (1991), in auctions with four or seven bidders, with $\epsilon = \$30$ and cash balances of \$4.50 (which forty-eight out of the fifty bidders always had), unilateral deviations from the RNNE bid function were not profitable.

49. Hansen and Lott go further, however, arguing that even if there were sufficiently large starting cash balances to cover all potential losses, there might still be a limited-liability problem in the initial auction periods due to multiperiod effects. This part of their argument does not apply to the Kagel and Levin design however (see Kagel and Levin 1991).

50. Lind and Plott report only one subject having to work off an \$8 loss, which he actually did. Obviously it may be difficult to enforce subjects working off losses. However, this can be dealt with by holding the participation fee in escrow and, when there are a number of different experiments taking place on a regular basis, the subject can be denied participation in these experiments until the debt is paid off.

51. To keep costs down the seller's auctions were conducted in francs as opposed to dollars. The conversion rate from francs to dollars reduced the cost of the experiment, but meant that losses due to departures from Nash behavior were substantially smaller in the selling experiment. Lind and Plott note that should otherwise inexplicable differences in behavior be observed across the two designs, these incentive differences would be an obvious line of research to pursue.

52. I conjecture that one factor underlying the higher bankruptcy rates reported in Cox and Smith results from the limited feedback bidders have regarding auction outcomes (only winning bids) versus the much more extensive feedback offered in Kagel and Levin (1986) and Lind and Plott (all bids and the corresponding signal values). The latter permits observational learning (determining the outcome of your bidding strategy conditional on having one of the higher, winning signal values) which helps bidders to adjust to the adverse selection problem underlying the winner's curse (Garvin and Kagel 1991; section II.D below).

53. The symmetric Nash equilibrium bidding function for second-price auctions was first reported in Matthews (1977). Milgrom and Weber (1982) extend Matthews' analysis to consider the effects of public information. Levin and Harstad (1986) showed that this function is the unique symmetric Nash equilibrium. Harstad (1991) extends the analysis to asymmetric Nash equilibria where the source of asymmetry is differences in bidders' risk preferences.
54. The intuition with respect to this strong result about the number of bidders is that as long as bidders rely on their private signal values to determine their bids (so that bid profiles are strictly increasing in x), then as n increases the expected value of the item conditional on winning decreases. And since this is a second-price auction, many of the strategic considerations that complicate the first-price auction are not present.
55. Milgrom and Weber (1982) develop this symmetric RNNE for the English clock auction. There are other symmetric equilibria as well, but all symmetric equilibria yield the same expected revenue (Bikhchandani and Riley 1991). This is the most plausible of the symmetric equilibria. Bikhchandani and Riley also discuss the existence of asymmetric RNNE for the English auction. The experiments reported here use an English clock procedure similar to the one described in section 1.2b. I assume $\underline{x} + \epsilon \leq x_L \leq \bar{x} - \epsilon$.
56. T-statistics here are 2.27 with $\epsilon = \$6$, significant at the 5% level (two-tailed t-test) and 1.98 with $\epsilon = \$12$, significant at the 10% level (two-tailed t-test).
57. The control group differs from Table 7.8. The latter auctions used senior undergraduates and night MBA students from the University of Houston. The data in Table 7.9 used primarily day MBA students from the University of Pittsburgh and some senior undergraduate economics majors.
58. For the asymmetric information auctions, $E(x_o \mid X = x_1)$ is calculated strictly on the basis of the outsider's signal values. With symmetry this provides a lower bound on the expected value of the item conditional on winning.
59. Actually, for $\epsilon = \$12$ the bid discount for the symmetric information auctions is significantly greater, at the 10% level. However, the control condition consists of two auction series, one of which exhibited unusually rapid adjustment to the winner's curse, with substantially higher average bid discounts than observed in any other symmetric information series. Further, the last twelve auctions in this session (with $\epsilon = \$12$) employed a dual market technique with subjects first bidding in a symmetric information market and then bidding in the asymmetric information market. In these auction periods the average difference in the bid discount across information conditions was substantially smaller than in Table 7.9 and was not significantly different from zero at conventional levels. This suggests that the differences reported in Table 7.9 resulted from a subject group effect rather than a treatment effect.
60. Expected profit here is $(b/200)(wb - 200)$, which for $b \leq 100$ is negative for all $w < 2$.
61. In the case of monetary incentives, subjects were responsible for losses out of their own pockets. The decision to play in this case was voluntary. Subjects were MBA students drawn from managerial economics classes, with the experiment serving as an introduction to decision making under uncertainty. Nineteen of 131 students chose *not* to participate with monetary incentives.
62. See Carroll, Bazerman, and Maury (1988) for a closely related study that uses verbal protocols in efforts to understand better the cognitive processes underlying buyers' decisions.
63. Samuelson and Bazerman (1985) had designs that addressed this issue as well. But they used one-shot trials with no allowance for learning.
64. Let p_A be the probability the seller accepts the bid b, let $v = c + V$, where c is the constant in the owner's valuation function, and let $V \in [0, R]$. Then the expected value of the item to the naive bidder is $E(G_N) = p_A\,[1.5(c + R/2) - b]$ and the expected gain to a rational bidder is $E(G_R) = p_A\,[1.5(c + b)/2 - b]$. The optimal naive bid maximizes $E(G_N)$ while the optimal rational bid maximizes $E(G_R)$. In the loser's curse $b_N < b_R$ and $E\,G \mid b_R) > E\,G \mid b_N)$. These inequalities are reversed with the winner's curse. See Holt and Sherman (1994) for further details.
65. Differences in the expected value of the naive bid and the rational bid are very small here—less than 2 cents. If anything this might be expected to bias the results in favor of the thrill of winning as the expected cost of deviating from the rational bid is quite small.

66. My analysis concentrates on the three pure common value auction series. The other three auction series combine common and private value elements, which complicates any evaluation of the winner's curse.
67. The average value of blind bid items in period 1 was 41.8 (with a standard error of the mean of 12.4) compared to an average value of 91.7 (with a standard error of the mean of 0.33) on those items that were not blind bid.
68. Average losses were 10¢ per item in the first ten periods compared to 4¢ per item from period eleven on. Excluding the positive profits made on one very high valued item that was blind bid out of seller's ignorance, or seller's error, mean profits averaged −6.6¢ with a standard error of the mean of 2.72, for period 11 on, so that buyers' profits on these blind bid items were significantly below zero. Including this one outlier, losses averaged −3.6¢ with a standard error of the mean of 4.0, which is not significantly below zero. Detailed reporting of the data in the *Rand Journal* facilitated these calculations. The *Journal* and the authors are to be applauded for this.
69. Information about product quality was delivered privately to buyers after their purchases. Public delivery of this information would undoubtedly speed up the adjustment process.
70. Mean earnings for the three subjects who did not return were slightly higher than those that did.
71. A Z statistic comparing the first quartile with the second quartile shows a statistically significant reduction in the proportion of auctions with losses from 64% to 41% ($Z = 2.24$, $p < .05$).
72. In the two buyer sessions reported in Table 2.7 there is a noticeable and unexplained increase in average predicted profits from playing the RNNE in the second quartile. This confounds any explanation of increases in actual profits in terms of individual learning.
73. A chi-square test indicates that this difference is significant at better than the .01 level.
74. The relevance of this market learning outside of laboratory markets is problematic. Entry conditions were closed in the experiment—subjects were recruited back exclusively from those who had participated in earlier auction market sessions. In field settings, open entry conditions permit new, inexperienced bidders (firms) to join the bidding. And it is not clear if these new players must undergo their own self-selection process and, if so, what the market impact would be since they would presumably make up a minority of active bidders.
75. Lind and Plott contain auctions with large monetary losses (the high price auctions) and with small losses (the offer auctions, which use an artificial currency, see Table 7.7). Their results are not sufficiently clear cut to reach a firm conclusion on this point. The threat of bankruptcy might also carry with it some embarrassment which would heighten subjects' sensitivity to losses.
76. Lind and Plott reported all bids and signal values to subjects as well.
77. These auctions were carried out as a research project in an undergraduate class in experimental economics. The undergraduate students involved were Robert Van Winkle, Danielle Rondelez, and Matt Zander.
78. In series 1 the reserve price realization was announced in the initial auction periods that did not permit communication. In series 2 the reserve price realization was not announced in the first five periods.
79. Pooling the data across auction series, a t-test shows these differences to be significant at better than the .01 level.
80. Minimal rationality requires the winning bid to be less than or equal to $E[x_o]$. With the notable exception of one period in auction series 2, this minimum rationality requirement was satisfied.
81. Dyer et al. (1989b) address the issue of whether the executives may not have taken the experiment seriously.
82. This despite some potentially important differences in rate of return methodology and drainage lease samples between the two studies. Although differences in rate of return between nonneighbor drainage leases and wildcat leases are not statistically significant ($t = 0.80$), neither are the differences in rates of return between neighbors and nonneighbors on drainage leases ($t = 1.13$).

83. In the experiment, one of the executives jokingly inquired "Who is the architect associated with this job?"
84. It is my understanding that different oil companies specialize in different geological formations so that they can better apply accumulated past knowledge to interpret seismic records.
85. For publicly owned projects, there are laws explicitly recognizing the possibility of "arithmetic" errors in bids, permitting the contractor to withdraw his bid without penalty. What constitutes an arithmetic error is often loosely interpreted. In this case, the low bidder used a plumbing subcontractor's bid which the subcontractor withdrew. However, the subcontractor was unable to reach all general contractors, the original low bidder included, in time to adjust their bids. What is clear in the construction industry is that no one wants a builder, or subcontractor, working for them who is terribly unhappy with their bid, as this affects the speed, quality, and "headaches" associated with the construction. Owners requiring that such bids be lived up to are likely to suffer from a winner's curse of their own. In OCS bidding, where the winner's curse expresses itself in terms of an inflated bonus bid, there are no corresponding considerations since the bonus bid is a sunk cost.
86. The ability to withdraw bids without penalty opens up additional strategic possibilities. However, in field settings, owners and contractors are involved in a game with two-sided reputations, so that too frequent withdrawal of bids may result in being left off future invited bid lists.
87. There are fewer bidders on average for OCS leases (3.5 versus 7.5 in the construction data). Limiting analysis of the construction data to jobs with 4 bidders or less results in auctions with an average of 3.4 bidders and average money left on the table of 6.7% (with a standard error of 8.8%) (Dyer and Kagel 1992).
88. In a frequently cited study, Thiel (1988) models highway construction bidding as a pure common value auction completely ignoring the important private value elements underlying bids in this industry. He finds no evidence of a winner's curse which may result from either this specification error or the modeling errors identified in Levin and Smith (1991).
89. Models developed to date deal with auctions in which buyers and sellers have a single unit to trade. In contrast, much of the experimental literature deals with two-sided auctions in which buyers and sellers each have several units available for trade.
90. ZI traders achieve remarkably high efficiency levels in continuous double auctions. Gode and Sunder (1993) conclude that the high efficiency levels commonly reported in CDA's probably has more to do with the structure of the institution than any particular skill on traders' part. The ZI model provides a zero rationality benchmark against which to compare outcomes.
91. This is a special case of what Satterthwaite and Williams refer to as the k-th price auction—the case of $k = 1$.
92. ZI traders achieve average efficiencies of 29% in markets with two buyers and two sellers and 36% in markets with eight buyers and eight sellers. In contrast, for the continuous double auction lower bounds on efficiency are around 90%. The main difference between ZI traders under a clearing house mechanism and the CDA is that the clearing house is completely unforgiving with respect to bids and offers that fail to achieve mutually beneficial trades, while the ability to make repeated bids and offers forgives such mistakes in the CDA.

Bibliography

Akerlof, G. 1970. The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics* 89:488–500.

Arrow, K. J. 1970. *Essays in the theory of risk bearing*. Amsterdam: North Holland.

Ashenfelter, O., and D. Genesore. 1992. Testing for price anomalies in real estate auctions. *American Economic Review: Papers and Proceedings* 82:501–5.

Ball, S. B. 1991. Experimental evidence on the winner's curse in negotiations. Diss., Evanston, IL: Northwestern University.