Reinhard Selten (Ed.)

Game Equilibrium Models II

Methods, Morals, and Markets

With Contributions by
D. Abreu, W. Albers, K. Binmore, R. Gardner
W. Güth, R. M. Harstad, H. Kliemt, W. Leininger
A. Okada, E. Ostrom, D. Pearce, L. Phlips
S. H. Schanuel, L. K. Simon, J. Sobel, G. Stephan
E. van Damme, J. W. Walker, F. Weissing, W. R. Zame

With 79 Figures

Springer-Verlag
Berlin Heidelberg New York
London Paris Tokyo
Hong Kong Barcelona
Budapest

A FRAMING EFFECT OBSERVED IN A MARKET GAME

Wulf Albers and Ronald M. Harstad

Introduction.

A wide variety of markets transacting via auctions share this characteristic: when a bidder selects his strategy, the value to him of the auctioned asset is uncertain. This characteristic lends the common–value auction model (Wilson [1977], Milgrom and Weber [1982]) much of its richness. It also makes the bidder's strategic problem more complex: laboratory observations of common–value auctions have suggested that subjects exhibit substantially more difficulty approaching equilibrium profit levels, and indeed avoiding losses, in common–value auctions (Kagel and Levin [1986], Kagel, Levin and Harstad [1988]) than in the more artificial setting of private–values auctions (where each bidder knows for sure the asset's value to him, cf. Kagel, Harstad and Levin [1987], Cox, Roberson and Smith [1982]).

In a canonical common–value auction, asset value is a random variable v. Each bidder privately observes a signal informative about v, typically these signals are independently and identically distributed given v. Based on this information, bidders compete under known auction rules. A symmetric equilibrium analysis of the auction has each bidder rationally take into account when selecting bids that, if he wins, all rivals will have estimated the asset to be worthless; the symmetric assessments yield this outcome because each rival observed a lower signal than the winner. Field observations of offshore oil lease auctions led Capen, Clapp and Campbell [1971] to suggest instead that bidders systematically fail to take into account the likelihood that the winner is the bidder who most overestimated the asset's value. This "winner's curse" is alleged to affect bidding and profitability in publishing (Dessauer [1981]), professional sports (Cassing and Douglas [1980]) and corporate takeovers (Roll [1986]). The laboratory experiments cited have found indications of the winner's curse to be widespread and persistent: more experience in a particular auction setting seems to be required for subjects to cope with this adverse selection problem than to cope with nearly any other game–theoretic issue presented in controlled laboratory environments.

We report here on a series of oral or "English" auction experiments designed to allow observation of the key characteristic of standard common–value auctions in a somewhat more simplified setting. The simplification arises principally in making subjects' information about asset value unconditionally independent. We conduct several replications of the following auction. Each of five bidders privately observes his "draw", a random variable drawn from a uniform distribution on the set {0, 5, 10, ..., 195, 200}. With this information, the five bid for an asset worth precisely the mean of the five draws. The draws are not signals, in that the asset value v is determined (nonstochastically) from the aggregate information contained in the draws, not vice versa. Presumably, the underlying statistical properties of this auction are more readily comprehended than in the canonical common–value model; the notions of a "low draw" or a "high draw" can be given unconditional meanings, and straightforward numerical interpretations.

The idea behind this approach was to reduce the internal variance of the model: the value of the asset does not contain any additional error term, and the sum of the information of all subjects together gives the correct value of the asset. Moreover, in symmetric equilibrium, at the end of every auction the asset value is precisely known to the winner. It can be calculated from rivals' dropout prices and his own draw.

As described in Section 2, the auction was conducted by having an auctioneer call out prices in ascending order, with subjects announcing when they wished to drop out, i. e., to cease competing. The asset was awarded to the last remaining competitor at the price at which his last rival dropped out of the bidding. The symmetric equilibrium model of this auction does have the bidder with the highest draw winning, and all bidders taking proper Bayesian account of this feature throughout the bidding. Bidders may fail to incorporate such Bayesian calculations, so a winner's curse hypothesis offers an alternative prediction of behavior; given the dynamic nature of English auctions, two formulations of the winner's curse are specified in Section 3.

Additionally, a treatment is incorporated that allows observing a pure framing effect in the market. The term framing effect comes from psychological studies of individual decision making under uncertainty, which suggest that two logically equivalent problems can lead to systematically different evaluations and decisions when questions are framed in different ways (Tversky and Kahneman [1981]). Economists have questioned whether markets as institutions discipline agents so as to eliminate or seriously constrain such irrational behavior (see, e. g., Knez and Smith [1987]). These experiments provide another vantage point on this issue, by observing whether a framing effect matters to market outcomes and to individual behavior in markets, in a setting presenting market participants with a task which may be at or beyond their rational capabilities.

The framing effect is introduced as follows. Experiments 1–5 were conducted as described above, under what we label the mean condition. Experiments 6–9 were conducted under the sum condition: each of five bidders privately observes his draw from a uniform distribution on the set {0, 1, 2, ..., 39, 40}. With this information, the five bid for an asset worth precisely the sum of the five draws. Notice the transformation of one condition into the other is completely transparent (multiplication of draws by 5, and mean instead of sum), and ought not affect rational behavior. Section 4 describes our observations as to whether it did.

1 Theoretical Predictions

Except where we explicitly consider differences, all discussion of theoretical predictions and of observations will be couched in terms of the mean condition draws. Thus, a bid associated with a draw of 35 should take into account that 35 is a relatively low draw (out of 200, corresponding to a sum–condition draw of 7 out of 40).

1.1 Symmetric Equilibrium

The underlying symmetry of the auction leads a game–theoretic analysis to predict a symmetric equilibrium solution. Here symmetry means that any two bidders with the same draw, observing the same history of dropout prices so far, will plan to continue competing up to exactly the same price. Equilibrium means that, given the behavior of rival bidders, no bidder can gain by a unilateral change of strategy. Techniques in Levin and Harstad [1986] (including showing that symmetric equilibrium must prescribe monotonicity, i. e. higher bidding at higher draws) can readily be adapted to the current auctions to show that symmetric equilibrium is unique. [1])

Symmetry and monotonicity imply that a bidder will be outbid if any rival has a higher draw. If a bidder is outbid, he does not care how high he bids. So the game–theoretic model has each bidder assuming that no rival observed a higher draw, an assumption that will (in equilibrium) be harmless when incorrect. This assumption guides plans as to how

[1]) It is well known that there exist asymetric equlibria which are degenerate, in the sense that one bidder always wins. Bikhchandani and Riley [1989] show that the symetric equilibrium is the only nondegenerate equilibrium, under a rather opaque condition that is satisfied in this model. Neither our observations or those of Harstad, Kagel and Levin [1989] lend any support to the degenerate equilibria.

far to compete before dropping out, with the effect that each bidder incorporates from the beginning information that would be available to him if he were to be the winning bidder. If a rival has a lower draw than yours, he will drop out sooner than you plan to drop out. So in essence, each bidder in equilibrium assumes that any rival still competing observed the same draw the bidder himself observed; this assumption is updated whenever a rival drops out, revealing that the rival actually had a lower draw.

It clarifies presentation and discussion to distinguish two rank orders. Let p_1, p_2, p_3, p_4 denote the nondecreasing sequence of prices at which a 1st, 2nd, 3rd, and eventually 4th bidder ceases competing ("dropout prices"). Thus, the market price paid by the winner is p_4, and the winner's profit is $v-p_4$. Let x_1, x_2, x_3, x_4, x_5 denote the draws observed by the bidders when ordered to correspond to the order in which they cease competing. That is, x_1 is the draw seen by the first bidder to drop out. Under symmetric monotonic behavior, x_1 is also the lowest of the five draws; otherwise it need not be. Let $x(1), x(2), x(3), x(4), x(5)$ represent the draws ranked in ascending order. Thus, $x(5)$ is the highest draw, while x_5 is the draw observed by the winning bidder.

The basic incentive structure of English auction rules is price-taking: your decision regarding how long to continue competing will determine whether you win, but the price you pay if you win will be determined by your rivals' decisions about how long to compete. The assumptions mentioned lead the bidder with the lowest draw to assume each rival shares his draw, thus v equals his draw, so $p^*_1 = x_1$. (Throughout, an asterisk will designate an equilibrium prediction. Recall that this prediction translates into a sum condition prediction $p^*_1 = 5x_1$.)

Note the implication of symmetry: each bidder initially plans to continue competing until the price reaches his draw, so long as no other bidder has dropped out. When they learn of p_1, each of the 4 remaining bidders re-evaluates his plans, estimating v at $0.2p_1 + 0.8x$ for the x he observed. Each plans to continue in the bidding up to this price, so $p^*2 = 0.2p_1 + 0.8x_2$. Once again, p_2 reveals to each of 3 remaining bidders that they must re-estimate v. Table 1 provides formulas and illustrates this logic for the mean draws and sets of draws used in the questionnaire.

While v is viewed as uncertain by each individual bidder, in equilibrium, for the winner this uncertainty is resolved as soon as his last rival drops out. At this point, v is inferentially nonstochastic, equal to the average of p_4 and the price at which the winner was planning to drop out $(0.1p_1+p_2/6+p_3/3+0.4x_5)$. Equilibrium profit is thus $0.2(x_5-x_4)$, which also becomes nonstochastic when p_4 becomes known. Ex ante, for x drawn uniform on $\{0, 5, ..., 200\}$, average equilibrium profit is 20/3.

1.2 Winner's Curse Models

The winner's curse phenomenon has not been given as precise a definition as equilibrium has. The fundamental distinction is that bidders falling prey to the winner's curse select bids under a presumption that there is no particular rank–order relationship between their draw and rivals' draws. In the auctions observed here, a bidder who recognizes that his draw of 20, say, is unlikely to be a median draw may nonetheless bid under a dangerous assumption in line with the winner's curse. That is, he may rely upon the independence underlying his draw and rivals' draws, assuming that he can bid as if rivals' draws are unrelated to his. A specific formulation of this view is what we label

> the **unresponsive winner's curse model**: assuming that rivals' draws each average 100, are independent of a bidder's draw, a bidder with draw x does bid up to 0.2x + 0.8*100.

If all bidders were to adopt this strategy, they would be behaving symmetrically, and the winner would be "cursed" to discover that his rivals' draws are related to his draw, which is an upper bound on their draws. (Even if four rivals' draws will average 100, the four lowest of five draws will average less.) The unresponsive winner's curse model has later dropout prices unrelated to earlier dropout prices; earlier dropout prices do contain the complete information about their draws (the bidder's draw is 5 times as far from 100 as his dropout price), but this information is hypothesized not to be used in this model. The unresponsive winner's curse model predicts an overall average profit of –20/3.

A bidder can begin the auction naively assuming there is no relationship between his draw and others, and still see advantages to behaving differently when rivals drop out early than when they compete longer. A version of a winner's curse effect remains if he assumes draws of rivals still competing are unrelated to his own draw. This consideration gives

> the **responsive winner's curse model**: (Counterfactually) rivals who have already dropped out are assumed to have dropped out at their draws (i. e., using p(x) = x), and rivals remaining are assumed to have draws averaging 100.

Table 1 provides the resulting predictions for mean draws and for sets of draws in the questionnaire. When a dropout price below 100 is observed, this model lessens the winner's curse by assuming an overestimate of the dropper's draw, but less of an overestimate than 100 would be. It still leaves bidders in serious jeopardy whenever x<100, as remaining

Table 1: Drop-out Prices in Equilibrium Bidding , and Responsive Winner's Curse Model for Three Examples of Auctions (general formulas at the bottom)

	Equilibrium Model: rank of dropout: first	second	third	fourth	Responsive Winner's Curse Model: rank of dropout: first	second	third	fourth
	dropout prices of bidders who dropped out before:							
	—	p1=33.33	p1=33.33 p2=60.00	p1=33.33 p2=60.00 p3=80.00	—	p1=86.67	p1=86.67 p2=90.67	p1=86.67 p2=90.67 p3=95.46
draw:	dropout price (if next to drop out):				dropout price (if next to drop out):			
x= 33.33	p1= 33.33*				p1= 86.67*			
x= 66.67	p1= 66.67	p2= 60.00*			p1= 93.33	p2= 90.67*		
x=100.00	p1=100.00	p2= 86.67	p3= 80.00*		p1=100.00	p2= 97.33	p3= 95.46*	
x=133.33	p1=133.33	p2=113.33	p3=100.00	p4= 93.33*	p1=106.67	p2=104.00	p3=102.13	p4=101.27*
x=166.67	p1=166.67	p2=140.00	p3=120.00	p4=106.67	p1=113.33	p2=110.67	p3=108.90	p4=107.89
	dropout prices of bidders who dropped out before:							
	—	p1=20	p1=20 p2=36	p1=20 p2=36 p3=48	—	p1=84	p1=84. p2=84.8	p1=84 p2=84.8 p3=85.76
draw:	dropout price (if next to drop out):				dropout price (if next to dropout):			
x= 20	p1= 20*				p1= 84*			
x= 40	p1= 40	p2= 36*			p1= 88	p2= 84.8*		
x= 60	p1= 60	p2= 52	p3= 48*		p1= 92	p2= 88.8	p3= 85.76*	
x= 80	p1= 80	p2= 68	p3= 60	p4= 56*	p1= 96	p2= 92.8	p3= 89.76	p4= 86.91*
x=100	p1=100	p2= 84	p3= 72	p4= 64	p1=100	p2= 96.8	p3= 93.76	p4= 90.91
	dropout prices of bidders who dropped out before:							
	—	p1=120	p1=120 p2=136	p1=120 p2=136 p3=148	—	p1=104	p1=104 p2=108.8	p1=104 p2=108.8 p3=114.56
draw:	dropout price (if next to dropout):				dropout price (if next to drop out):			
x=120	p1=120*				p1=104*			
x=140	p1=140	p2=136*			p1=108	p2=108.8*		
x=160	p1=160	p2=152	p3=148*		p1=112	p2=112.8	p3=114.56*	
x=180	p1=180	p2=168	p3=160	p4=156*	p1=116	p2=116.8	p3=118.56	p4=121.47*
x=200	p1=200	p2=184	p3=172	p4=164	p1=120	p2=120.8	p3=122.56	p4=125.47
drop-out price (as function of own signal x, others' draws xi / others' dropout prices pi) by rank of dropout:								
first	p1= x				p1= .2x + .8			
second	p2=.2x1+.8x =.2p1+.8x				p2= .2p1 + .2x + .6			
third	p3=.2x1+.2x2+.6x =.15p1+.25p2+.6x				p3= .2p1 + .2p2 + .2x + .4			
fourth	p4=.2x1+.2x2+.3x3+.4x=p1/10+p2/6+p3/3+.4x				p4= .2p1 + .2p2 + .2p3 + .2x + .2			

* Dropout prices of bidders who actually drop out next (if the draws are as given in the left column) are marked by an asterix.

rivals' draws are overestimated when it matters (i. e., when the bidder wins). For x>100, however, once two or three rivals have dropped out, 100 will be on average an underestimate of rivals' draws, so this model can readily observe profits when v>100. Overall, it predicts an average profit of −1.23.

The two formulations used to yield predictions of a winner's curse story are unaffected by the framing effect of switching between mean and sum conditions. Psychologists' arguments for the relevance of framing effects are based upon a decision process where subjects adopt an anchor and then systematically underadjust their estimation away from that anchor. The way a question is framed may alter the anchor selected. In this context, it is not clear to us how differential predictions of the winner's curse may be built by hypothesising that the mean and sum conditions lead to subjects adopting different anchors. We do not rule out the possibility; we simply are unsure how to model it.

2 The Experiment

The subjects were recruted from a microeconomics course at the University of Bielefeld. They were mainly business administration students. The experiment was run in the same room as the course, but the participation was voluntarily. 45 students participated, i. e. about one third of the students enrolled. The subjects were distributed to 9 different tables in a way that students who were acquainted with each other sat at different tables.

The instructions were supported by a one page leaflet containing all essential information. (The modifications for the sum condition are given in brackets.)

> **The Problem:** An asset will be auctioned among 5 people in an English auction. The Value of the asset is the mean (sum) of 5 individual values each of which is known to one of the 5 persons. – The 5 individual values are independent draws from the interval 0 to 200 (0 to 40). Every individual value can adopt the values 0,5,10,15,...,195,200 (0,1,2,3,.. ..,39,40), and each of these values has the same probability. – Example: individual values 135,75,110,170,45 (17,25,22,34,9) give a mean (sum) of 107, which is the value of the asset.
>
> **The English auction:** An auctioneer counts aloud the numbers 0,1,2,3,... to 200. A bidder does not act as long he is willing to pay the announced number as the price of the asset. As soon as the auctioneer says a number which is higher than a bidder (for instance bidder A) is willing to pay, the bidder says "A off". This means that Bidder A is not bidding any more from then on. – The auction is done, as soon as

only one bidder is still bidding. He is the buyer. The number, at which the last of the other bidders said "X off", is the price of the asset. (profit = value – price)

Remarks to the rules of the auction: The decision "X off" is irrevocable. – If several bidders say "X off" at the same price, that one receives the asset who said "X off" last. (This decision is made by the auctioneer, if necessary, by flipping a coin.)

Remark concerning the execution: The individual values have been taken in advance as independent draws, and assigned to a separate list for every player. Do carefully pay attention, to pick up your individual values for the correct rounds. (It seems to be helpful, to erase each value after transferring it to the auction form.)

Aim (German: Zielsetzung): Experienced bidders can conclude from their individual number and the bidding behavior of the others on the mean (sum), and can behave accordingly.

These Instructions were read aloud to all subjects simultaneosly. – Before separating the groups and reading the instructions the experimentor shortly motivated the two different framing conditions by two different cover stories:

Mean: Five firms are interested in the rights to strike oil from a certain area. They have different estimates of the value of these rights, given by the five best experts in this field. Every firm does only know the estimate of one expert. The correct value is precisely the mean of the expert opinions.

Sum: Five firms are interested in getting the rights to strike oil from a certain area. The area can be subdivided into five subareas. In each of these subareas one of the firms drilled down a testwell, giving the precise value of this subarea. The result of this test is not known to the other firms. The value of the total area is precisely the sum of the values of the subareas. The area is auctioned in one piece.

The groups ran between 48 and 60 auctions in a time of around 4 hours including instructions. Five groups were run under the mean, four under the sum condition.

All subjects (and the supervisors = auctioneers) had identical forms in which they filled in the dropout prices of all players, the true value of the asset and the calculated profit for each single auction, so that it can be taken for sure that every player really perceived these data. The auctioneer had a reservation price which was 20 units below the value of the asset in the respective round. When the last dropout price did not meet the reservation price, the asset was not sold and the subjects were not informed about the exact value in this round. Otherwise the true value of the asset was given by the auctioneer after every auction. Information about the individual values was not given at any time.

After every sixth auction, before the value of the asset was reported, the subjects were asked to guess the value.

After every twelfth auction, a separate questionnaire was distributed asking for the dropout prices in different situations:

> Question 1: You are bidder 1. Your individual value is x_1 (=20 or 120). At which number p_1 do you say "X off", if no one else stops bidding before you?
>
> x_1 = 20 —> p_1 = ...
> x_1 = 120 —> p_1 = ...
>
> Question 2: You are bidder 2. Your individual value is x_2 (=40 or 140). At which number p_2 do you say "X off", if no one else stops bidding before you? (Fill in p_1 from above.)
>
> p_1 = ... , x_2 = 40 —> p_2 = ...
> p_1 = ... , x_2 = 140 —> p_2 = ...

Questions 3 and 4 gave the same type of question for the fourth and the third bidder who dropped out, where the individual values were (60 or 160) for the third and (80 or 180) for the fourth. Thereby two increasing sequences of individual values were asked, namely 20,40,60,80 and 120,140,160,180, with the second 100 points higher than the first. Each of the corresponding questions was based on "observed" dropout prices p_1,p_2,.. of the dropout responds of the same subject in the preceding questions, and on the individual information of the respective question (x_1 or x_2 or..). In this framework the players had the opportunity to give a complete set of decisions in all positions. It might even be suggested that he was motivated to decode the preceding dropout signals, in order to reach a reasonable result. (In advance we feared that the decoding might be done just by looking at the individual values of the preceding questions. But the subjects did apparently not behave like that.)

The value of each point was 1 DM (over 1/2 Dollar). This meant that if they played the equilibrium dropout prices, they could win 300 DM per table. If they cooperated, and did all drop out at low prices, they could have even made a profit up to 100 DM per game. To prevent such cooperation the subjects were told that every asset had an announced reservation price, below which it would not be sold. (The reservation prices were 20 points below the values, but this was not known to the subjects.)

The subjects were told that they did have an initial endowment of 60 points, and that in similar experiments the most successful players received total profits of about 150 DM (about 75 Dollars) to 200 DM (about 100 Dollars). (Anyway the subjects did not suspect in advance that they might make losses in the auctions.) The players were instructed that no one would lose money in the experiment, but that they should contact the experimentor, if high losses in the preceding sessions caused, that they did not feel motivated any more. (One subject did so, and we raised his aggregated profit at that moment from –300 up to –50.) During the experiment, facing the high losses by the winners curse, the players were told that other players were also making losses, and that it might be that we would give some additional endowment to all of them. – The final payoffs of the experiment were based on an initial endowment of DM 60. The most successful players received DM 175, 121, 111, 105, and 84. The sum of all payoffs was 1215 DM (about 610 Dollars).

Table 2: Mean Nash–Equilibrium Profit, and Mean Observed Profit *)

Exp.	# aucts	mean NE profit all auctions	mean observed profit all auctions	mean observed profit auctions 1–48	mean observed profit auctions 37–48
1M	72	6.97 (0.79)	-3.74 (2.24)	-2.73 (2.71)	-9.92 (5.86)
2M	60	6.53 (0.72)	-2.46 (1.91)	-0.43 (2.04)	-2.25 (3.90)
3M	60	6.48 (0.63)	-8.20 (2.65)	-9.44 (2.74)	-13.33 (3.93)
4M	60	7.18 (0.73)	-6.28 (1.96)	-4.94 (2.28)	-4.00 (3.84)
5M	54	5.63 (0.75)	0.19 (2.79)	-0.13 (3.06)	6.25 (5.34)
1S	60	6.75 (0.85)	-1.98 (2.40)	-2.25 (2.49)	-1.17 (6.21)
2S	48	6.58 (0.75)	-3.50 (2.02)	-3.70 (2.03)	1.67 (3.11)
3S	56	6.85 (0.64)	-7.42 (2.58)	-7.10 (2.80)	-5.25 (4.89)
4S	48	7.40 (0.83)	-3.69 (2.39)	-3.69 (2.39)	-7.42 (6.23)
1M–4M	252	6.80 (0.36)	-5.13 (1.12)	-4.40 (1.25)	-7.38 (2.32)
1S–4S	212	6.85 (0.39)	-4.09 (1.20)	-4.19 (1.23)	-3.04 (2.68)

*) The terms in brackets give the respective standard deviations

3 Results

3.1 Observations of Market Prices

Aggregate outcomes are compiled in Table 2. Experiments are identified by a number and letter M for mean condition or S for sum condition. Experiments with the same number had identical realisations for all draws. Rows labelled 1M–4M and 1S–4S combine observations for the indicated experiments. The number of auctions completed in each experiment is shown in column 2. Equilibrium predicted mean profits are shown in column 3. Mean observed profits are reported in columns 4, 5 and 6, for all auctions, for auctions 1–48, and for auctions 37–48 (the latest dozen completed in all experiments). Each mean profit entry is followed by its standard error in parentheses. Table 3 reports two frequencies of interest: the frequency with which nonnegative profit levels were observed, in column 2, and the frequency with which all five bidders were still competing when the price exceeded the value of the asset, in column 3.

Figures 1–4 diagram observed and predicted profit for the sets of matching experiments. Time progresses along the horizontal axis. The profit that would have been observed in symmetric equilibrium in each auction is indicated as a dollar sign; the actual profit observed in experiment 1M is indicated as an open square, and that observed in experiment 1S as a plus sign.

The overall impression is clearly an impact of some version of the winner's curse which is widespread and persistent. Average profit over the last 12 auctions only reaches a nonnegative level in experiments 5M and 2S. Both levels and frequencies of losses are not significantly different between the mean condition and the sum condition (mean profit for experiments 1M–5M is –4.18). The following null hypotheses cannot be rejected by either parametric differences in means tests or two–sample nonparametric rank tests:

a. Levels of profits are drawn from the same distribution in mean and sum conditions.

b. Signs of profits are drawn from the same distribution in mean and sum conditions.

c. Levels of profits are drawn from the same distribution in auctions 4–27 as in either auctions 37–48 or the last dozen auctions.

Figures 5–9 display profit as a function of value (centered on its mean of 100), both as scatter plots and as a regression line. Figures for experiments not shown are similar, as the tendency to observe losses whenever value was below its mean level of 100 and typically profits when value was above its mean level is robust across experiments and experience. Market prices show a mean reversion tendency sufficiently strong to generate this profit pattern. To understand this tendency, it is necessary to examine observed strategies as indicated by dropout prices.

3.2 Observed Dropout Prices

As indicated above, the symmetric equilibrium predictions begin with the subject with the lowest draw $x(1)$ dropping out when the price reaches his draw (or 5 times his draw in the sum condition). 44 of the 45 subjects continued to compete past this predicted level on at least 90% of the occasions when they held the lowest draw.

The excessively high levels of p_1, the first dropout price, are significantly different for the two treatment conditions. While a hypothesis of homogeneous behavior when dropping out first can be rejected in all experiments (via f–tests on fixed–effects regressions, $p < 0.1$), bidders in the sum condition exhibit significantly higher ($p < 0.001$) first dropout prices, both in nonparametric paired comparisons and in fixed–effects regressions of p_1 on the draw of the bidder dropping out first. Combining within conditions, OLS regressions of p_1 on the draw held yield $p_1 = 56.9 + 0.33$ draw, for the mean condition, and $p_1 = 73.4 + 0.25$ draw (that is, + 1.25 times an [unadjusted] draw ranging from 0–40), for the sum condition. Recall that the equilibrium model predicts a zero intercept and a slope of 1; either winner's curse model predicts a 0.2 slope and an intercept of 80.

Table 3: Auctions with Positive Profits, and Auctions where the First Dropout Price p_1, was Greater that the Value of the Asset (Proportions within the First 48 Auctions

Exp.	% of auctions with profit>0	% of auctions with p1>value	Exp.	% of auctions with profit>0	% of auctions with p1>value
1M	50.0	22.9	1S	43.8	29.2
2M	46.8	0.0	2S	43.8	29.2
3M	33.3	20.8	3S	40.9	29.5
4M	45.8	22.9	4S	45.8	31.2
5M	50.0	16.7			
1M-4M	43.9	16.8	1S-4S	42.7	29.2

This difference also shows up in columns 3 and 6 of Table 3, where the first drop out has already ensured a loss in 17% of the cases for mean condition experiments, but 29% of the cases for sum condition experiments. It is not easy to provide a thorough explanation for this difference; clearly subjects fall prey to the framing effect in some manner. Possibly the mean and sum conditions lead to anchoring and adjustment decisions with different anchors (as if first dropouts in the mean condition were calculating how far above their draw to drop out, while first dropouts in the sum condition were calculating how far below 100 to drop out).

Fixed–effects regressions of p_1 on the draw held by the first person to drop out, with or without time trends or lagged profit, explain between 50% and 60% of the variance in p_1 (except for 80% in 4M and only 37% in 3S). Between the small coefficients on the draw, and the remaining 40–60% unexplained variance, there is little opportunity to infer the draw from knowledge of p_1. This difficulty by itself must bear considerable responsibility for losses occurring in 57% of the auctions, and explain why mean condition auctions cannot take greater advantage of relatively less overbidding by the first bidder to drop out.

Behavior of the second, third and fourth bidders to drop out appears as if they are somehow disciplined to keep the impact of the framing effect upon p_1 from influencing p_4, the market price. Fixed–effects regressions of p_2 on x_2 are notably less different across the mean/sum treatment than the comparable regressions of p_1 on x_1: only two of the four paired comparisons of p_2 regressions exhibit average absolute values of subject dummy variable coefficients that are different across mean/sum conditions at 0.1 significance. Corresponding fixed–effects regressions of p_3 on x_3 and p_4 on x_4 or on x_5 (no 5th dropout price is ever revealed) show no differences at all across the framing effect.

When one or two bidders have dropped out, a suspicion that they have remained competing at prices above their draws should lead bidders still competing to drop out earlier than they otherwise would. However, this tendency should show up as an intercept adjustment in a linear bid function—it should not reduce the slope below 0.2 (the prediction of the equilibrium and both winner's curse models), as a higher draw still impacts positively on value. To examine this, for each experiment, p_3 (and separately p_4) was regressed on earlier dropout prices, a time trend, and x_3 (separately, x_4). Individual intercept coefficients for each subject were estimated, as usual for fixed effects estimates; individual slope coefficients were also estimated. Of the 45 slope coefficients on x_3, no estimate exceeded 0.15. Of the x_4 coefficient estimates, one was 0.08 and significantly positive ($p < 0.2$, 2–tailed), another was (counterintuitively) –0.10 and significantly negative; the remaining 43 slope estimates were below 0.12 and insignificant. The tendency of p_1 to be well in excess of x_1 clearly left little room for maneuvering to avoid losses when $v < 100$.

Some of the explanation for profits observed to increase with v for $v > 100$ relates to p_3 and p_4 behavior which provides virtually no information about x_3 and x_4.

Without depending upon the distributional assumptions underlying fixed-effects regressions, it is possible to see how well various rules of thumb organize the data. Across all sum condition auctions, the unresponsive winner's curse model (drop out at $.2x + .8*100$) has an average error of 0.16 in predicting p_2. The responsive winner's curse (drop out at $.2x + .2p_1 + .6*100$) does a much better job of organizing p_2 observations in mean condition auctions.

Notice that every subject regularly observes outcomes where $v < p_1$, outcomes where v is so slightly in excess of p_1 as to make the likelihood that $p_1 > x_1$ prohibitively high, and outcomes where he himself dropped out first at a price above his draw. So each should be aware of the need to consider p_1 as an upwardly biased estimator of x_1, and possibly an even more biased estimator of $x(1)$. The evidence suggests they do not attempt this crucial task: several rules of thumb which would take into account $x_1 < p_1$ when determining p_2 do not organize the data as well as the two rules of thumb just mentioned. Yet the unresponsive winner's curse sets p_2 at the asset's expected value given x_2 and assuming the other four draws (including x_1) average 100, taking no account of any information inferable from p_1. The responsive winner's curse sets p_2 at the asset's expected value given x_2, assuming (counterfactually) $p_1 = x_1$ and assuming the other four draws average 100.

While the unresponsive winner's curse predicts p_2 on average 0.5 points higher than observed in sum condition auctions, and 2.3 points lower than observed in mean condition auctions, it is an extremely high variance predictor in either. Variance accounted for (via measuring root prediction squared error), among simple rules of thumb, the best we have found at predicting p_3, in either treatment, is that the bidder drops out as soon as legally permitted after the previous dropout. The same rule of thumb is the best simple predictor of p_4 observations in both treatments. (Both the unresponsive and the responsive winner's curse fare very badly as predictors of p_4 behavior.) In fact, respectively in mean and sum conditions, p_4 falls within a point of tying p_3 with 49.7% and 60.1% frequency; of these ties, 38.9% and 33.6% happen to be cases where $x_4 < x_3$, and 50.5% and 54.9% happen to yield losses. Allowing for overlap, 68.4% and 69% of ties end up being cases where $x_4 > x_3$ and a profit resulted. Notice that the most workable simple predictors of p_3 and p_4 take no account even of the information the bidder dropping out had: his own draw.

There are some indications that the ties observed may have related to rational attempts to cope with not entirely hopeful situations (cf. Table 4). Consider the second bidder to drop out: in mean condition experiments he ties the first dropout 17 % of the time. These ties were cases where $x_2-x_1=-7$ on average, while cases where $p_2>p_1$, x_2-x_1 averaged 21. In the sum condition, where p_1 reflected substantially more overbidding, p_2 tied p_1 with 33 % frequency, on average x_2-x_1 was 10 for ties, 35 otherwise. These distinctions for cases of ties are essentially duplicated for p_3. The last bidder to drop is substantially more likely to tie in the mean condition (41 %) and somewhat more likely in the sum condition (53 %). For ties, x_4-x_3 averaged 13 (mean), 29 (sum). When the fourth bidder competed longer, x_4-x_3 averaged 12 (mean), 53 (sum). Overall the ties were occurring when the tying bidders signal was low relative to the current price, which is consistent with essentially more frequent ties in the sum condition.

Table 4: Mean Draws, Mean Drop–out Prices, Mean Values for Different Drop–out Ranks (First 48 Auctions)

		value = mean of individual draws					value = sum of individual draws				
rank of dropout		% ties	mean draw *)	mean dropout price *)	mean value o.asset	value./. dropout price	% ties	mean draw *)	mean dropout price *)	mean value o.asset	value./. dropout price
first		—	51	74	98	24	—	45	85	99	14
second	ties	17	61(-7)	92(—)	101	9	35	56(10)	94(—)	94	0
	no ties		68(21)	89(19)	97	8		79(35)	95(15)	102	7
	all		67(13)	89(15)	98	8		71(26)	94(10)	99	5
third	ties	29	69(-10)	99(—)	93	-6	46	77(6)	96(—)	93	-3
	no ties		99(38)	96(10)	100	4		105(32)	101(8)	104	3
	all		91(24)	97(8)	98	1		92(20)	99(5)	99	0
fourth	ties	41	103(13)	95(—)	90	-5	53	118(29)	99(—)	95	-4
	no ties		143(12)	107(9)	103	-4		143(53)	108(8)	103	-5
	all		127(12)	102(5)	98	-4		132(41)	103(4)	99	-4
winner			154(28)	—	98			153(21)	—	99	
			mean price payed:		102	-4		mean price payed:		103	-4

*) The terms in brackets give the mean additional draw (with respect to the draw of the player who dropped out before), and the mean additional dropout price (with respect to the same player).

Both regression analysis and comparative evaluation of rules of thumb, then, lead to the conclusion that the framing effect noticeably alters the behavior of the first bidder to drop out, but quicker following dropouts in sum conditions keep this effect from impacting on market prices, profits and p_3. There is little evidence that much valuable information can be gleaned from the prices at which the first couple bidders drop out, and much less evidence that bidders incorporate in their behavior in any systematic information from preceding dropout prices.

3.3 Information and Learning by Punishment

The idea that adequate bidding behavior is a matter of learning sounds reasonable, and can help to explain the observed results. Since the only player who can be punished for inadequate behavior is the winner of the asset, learning should – at least in the first phase of the learning process – mainly address the dropout price of the last player who drops out (p_4). If he drops out too early, the asset is left to a rival, at "too low" a price (that is, a price where the winning rival on average makes a nonnegligible profit, while of course the last bidder to drop out makes zero profit). – A basic adjustment of the intercept of the bidding function can then serve to keep overall mean losses reasonably small (nearly zero), but such simple adjustment would not take into account of information that may be available from the earlier dropout prices.

Of course, there are very simple rules, by which an overall profit of 0 can be reached. For instance if all bidders drop out at 100. But such a behavior would not be in equilibrium, since by dropping out later if own draws are above the mean and dropping out earlier at lower own draws gives additional profits. Moreover the dropout prices of the others can be used to get additional information about the asset. As we mentioned above, in equilibrium bidders do get the total information about the rivals' draws, but this was clearly not observed in the experimental behavior.

This raises the question, how much of the information about the individual draws was aggregated by the auction. It seems reasonable to take the dropout price p_4 (which is the price) as a conservative estimate of the aggregated information. This assumption can be motivated in two ways:

(1) Bidders see that they incur losses, when they drop out above asset value, and that they leave profits to others, if they drop out at prices below the value. This suggests that they – in the mean – should just drop out at the value.

(2) A more detailed analysis gives the result that this bidder (bidder 4), when deciding on his drop out price, has – explicitly or implicitly – (a) to make an estimate of rivals' draws, (b) to add his own draw, and (c) to estimate the last competing bidders' draw and add it. Although it is not unreasonable that such a behavioral idea guided the intentions of the players, the exact analysis gives that under normal conditions the bidder should learn to drop out earlier, namely at the price obtained by adding his own draw x_4 to (a) and (b). (Assume a monotonic

dropout behavior, i.e. that a bidder with a higher draw drops out at a higher price: he will be punished, if he drops out at a higher price, since if he then wins the asset, then his price is the higher one and the draw of the remaining bidder is overestimated. He also is punished for a lower dropout, since he afterwards can observe that the remaining bidder makes still good profits, a part of which he could have received had he won the auction.) – This means that, at the highest possible level of information, the dropout price of the price setting bidder can only reach $(x_1+x_2+x_3+x_4+x_4)/5$. I.e. p_4 is for (x_5-x_4) below the value $(x_1+x_2+x_3+x_4+x_5)/5$.)

Of course approach (2) is more correct, the remaining part of information, x_5-x_4, is not available to the subjects, and any estimate of this can theoretically only shift up the reaction curve by a constant. However, it cannot be excluded, that in a situation with personal contact part of this remaining information x_5-x_4 can even be transposed to the other (for instance by observing the trembling of fingers, etc.). This argument makes it reasonable to take p_4 as a slightly conservative indicator on how much of the information $x_1+x_1+x_3+x_4+x_5$ has been incorporated.

As can be inferred from figures 5–9 observed behavior was less informative about bidders' draws. Since profit is $(x_1+x_2+x_3+x_4+x_5)/5-p_4$, one minus the slope of the regression line in figures 5–9 is a conservative estimate of the amount of information about the draws which is aggregated in p_4. So the observed behavior on average only aggregated 50.79 % (in sum condition) of this information, leading bidders who were clearly seeking profit to the level shown in these figures. (The regression line in figures 5 and 6 intercept $v=100$ at about $p_4=104$.)

The intercept is a result of the overall learning behavior of the subjects over all games, avoiding expected losses, where the level of 104 shows that they did not yet learn adequately. This level has in the average been reached already after the first 12 rounds, and it seems that no more essential learning took place after that with respect to the mean level of losses.

Clearly, a bidder would also have reasons to deviate from his own most reasonable estimate into a more conservative dropout price. However this would – if done willingly – only effect the (slope for) high values, and not the unsuccessful behavior at low values v. The fact that the figure shows a linear shape of the in section suggests that a splitted behavior did not take place, and that the subjects everywhere used all the information they perceived. (Systematic underestimations or overestimations do only effect the intercept, not the slope.)

The data show that different groups performed essentially different in information aggregation. It seems surprising that the two groups on tables 2M and 2S (with identical draw data) performed best, and that the groups of the corresponding tables 3M and 3S (again with identical draw data) performed best. This slightly suggests that good information transferring can be motivated by the history in which signals are presented.

3.4 Bidders' Estimates of Value

Recall that every sixth auction involved an extra step: after the four dropout prices were known, the bidders were asked to estimate the value, before its announcement. Subjects appeared to take this request seriously. The dasta of table 5 show that the bidders, who dropped out first, second or third (on average) underestimated the value of the asset (after knowing all four dropout prices). Players who dropped out fourth, and "actively" set the price (not tying), on average underestimated the value by 3 points (mean), and 7 points (sum), and thereby implicitly estimated losses of 3 and 7 points. Winning bidders, however, overestimated the value by 8 points (mean), and 4 points (sum), and thereby implicitly estimated profits of the same amounts. The actual average losses were 5 points (mean), and 0 points (sum) in the respective rounds. Overall the data illustrate that on average the prrocedure selects just those players as winning, who do oversestimate the asset.

A question is, why the winning bidder was willing to bid higher than he thought the asset was worth? We initially suspected an answer comes from the frequently closely bunched occurrences of p_2, p_3 and p_4: in such cases, when (unknown to the bidders) the auctioneer was calling out a price equal to the asset's value, there may have been 3 or even 4 bidders still competing, all thinking at the moment that the value is higher. When a couple bidders drop out shortly thereafter, the price setter may have rationally viewed these bunched dropouts as suggesting a lower value than he thought before hearing such close dropouts. In fact the (implicitly) estimated losses of the 4th to dropout were on average for 7 points (mean), and 6 points (sum) higher in cases of ties than otherwise: Tying fourth dropouters are more cautios in evaluating the asset than those who do not tie. However, the winner on average overestimates the asset in both cases, in the sum condition even for the same amounts.

The price setter's estimates may become slightly more accurate in later auctions than earlier, but overall there is little sign of improvement in estimating accuracy as bidders gather experience. In view of the regression analyses which suggested that draws x_4 play virtually no role in determining p_4, it may be surprising that the winner regularly estimates the value about 10 points higher than the price setter, and this more or less matches the average of $x_5 - x_4$, the difference in their draws (9 points). This compares averages, however, of two sets of observations with extremely high standard errors. Of course, we have no corroborative evidence in the bidding: there is no way to know at how much higher a price the winner would have wanted to cease competing.

Table 5: Mean Draws, Mean Drop–out Prices, Mean Values, Mean Estimated Values for Different Drop–out Ranks (Every 6th of the First 48 Auctions)

		value = mean of individual draws					value = sum of individual draws				
rank of dropout		% ties	mean draw 1)	mean dropout price 1)	mean val.of asset2)	value./. dropout price 2)	% ties	mean draw 1)	mean dropout price 1)	mean val.of asset2)	value./. dropout price 2)
first		—	69	86	107(103)	21(17)	—	53	88	108(94)	20(6)
second	ties	17	50(-57)	103(—)	114(101)	11(-2)	42	63(6)	98(—)	101(92)	3(-6)
	no ties		79(18)	100(18)	106(98)	6(-2)		112(62)	102(21)	113(104)	11(2)
	all		74(5)	100(15)	107(98)	7(-2)		91(38)	100(12)	108(99)	8(-1)
third	ties	14	64(0)	98(—)	87(94)	-11(-4)	54	93(0)	99(—)	104(96)	5(-3)
	no ties		107(31)	108(7)	111(106)	3(-2)		104(15)	111(9)	112(114)	1(3)
	all		101(26)	106(6)	107(104)	1(-2)		98(7)	104(4)	108(104)	4(0)
fourth	ties	38	77(13)	104(—)	97(106)	-7(+2)	73	119(25)	103(—)	104(100)	1(-3)
	no ties		157(61)	116(9)	114(111)	-2(-5)		159(49)	119(11)	117(110)	-2(-9)
	all		127(26)	112(5)	107(109)	-5(-3)		129(31)	108(3)	108(103)	0(-5)
winner			166(40)	—	107(115)	(3)		167(38)	—	108(112)	(4)
			mean price	payed:	112	-5		mean price	payed:	108	0

1) In data–columns 2 and 3 the terms in brackets give the mean additional draw (with respect to the draw of the player who dropped out before), and the mean additional dropout price (with respect to the draw of the player who dropped out before).
2) In data–columns 4 and 5 the terms in brackets give the mean estimated value (by questionnaire), and the differene of this estimated value to the dropout price (i.e. the "estimated profit").

3.5 Responses to Questionnaires

In addition to providing estimates every six auctions, recall that the subjects were asked to fill out a questionnaire every twelfth auction. (Details are in the last part of Section 2.) The average responses to the questionnaire are presented in Table 6. Averages from mean and sum conditions are combined, as they are virtually identical. Overall, there is no clear pattern showing improvement in bidding on questionnaires as subjects acquire more experience (experiment 2M clearly improves in $q_1(20)$, and 1S improves in $q_4(180)$, but these are balanced by disimprovements in other experiments). The spread $q_4(80) - q_1(20)$ is systematically narrower than the spread $q_4(180) - q_1(120)$ in individual responses.

The questionnaire provided a means of observing hypothetical choices of the same subject in all four dropout positions of a low–draws and then a high–draws auction, just shifted by a constant of 100. It could also have served to suggest behavioral patterns for bidding, by pointing to the notion of bidders dropping out in the order of their draws. There is no evidence it actually served this purpose. Apparently subjects responded as we had hoped, not with separate calculations, but more with the same behavioral patterns they were using for the bidding.

Table 6: Average Responses to Questionnaires in a Low Draws and a High Draws Auction (Equilibrium and Responsive Winner s Curse Predictions in Parentheses)

rank of dropout	first		second		third		fourth	
questionnaires, low draws auction:								
own draw	x1	20	x2	40	x3	60	x4	80
avg dropout price	q1	72	q2	79	q3	84	q4	88
(equilibrium)	(p^1)	(20)	(p^2)	(36)	(p^3)	(48)	(p^4)	(56)
(resp.winner's curse)	(p^1)	(84)	(p^2)	(85)	(p^3)	(86)	(p^4)	(87)
questionnaires, high draws auction:								
own draw	x1	120	x2	140	x3	160	x4	180
avg dropout price	q1	110	q2	119	q3	126	q4	132
(equilibrium)	(p^1)	(120)	(p^2)	(136)	(p^3)	(148)	(p^4)	(156)
(resp.winner's curse)	(p^1)	(104)	(p^2)	(109)	(p^3)	(115)	(p^4)	(121)
observed behavior:								
avg own draw	x1	48	x2	69	x3	91	x4	129
avg dropout price	p1	79	p2	92	p3	98	p4	102
(equilibrium)	(p^1)	(33)	(p^2)	(60)	(p^3)	(80)	(p^4)	(93)
(resp.winner's curse)	(p^1)	(87)	(p^2)	(91)	(p^3)	(95)	(p^4)	(101)

If the pattern of higher dropout prices relating to higher draws is presumed to persist, then the value could be expected to fall in the range [56,80] and average 68 for the low–draws auction, and in the range [156,160] and average 158 for the high–draws auction. This would mean that most subjects respond to the low–draws questions with behavior that has 3 subjects still competing when the price has gone beyond the highest possible value of 80. They respond to the high–draws questions with behavior that does not even attempt to compete for profits (with $x_2 = 140$, inferring $x_1 = 120$, and x_3, x_4, $x_5 > 140$, the 2nd, 3rd and 4th bidder all know $v > 136$, yet all three drop out before that). With the high draws exceeding the low draws by 100, a value 20 points higher is ensured, so responses should be 20 points higher only if there is no reason to assume rivals' draws are higher. There is reason to assume this, and responses are more than 20 points higher in 90% of the cases. However, if rivals still competing are assumed to have as high a draw as yours, and higher draws are inferred for rivals who dropped out at much higher prices, responses 100 points higher are justified. The observed tendency to compete only 40 points higher shows an assumption that rivals have higher draws, but little confidence in the inferring from higher dropout prices. (The observed increases from low–draws to high–draws questions are larger than predicted by either winner's curse model.)

Responses are not only less responsive to the 100 point difference than would seem appropriate, but the responsiveness to higher draws (20 above the previous draw) is insufficient. It is well above (except for $q_4(80) - q_3(60) = 4$) the 4 points justified if you assume no rival's draw is higher, but averages roughly half the responsiveness warranted if you assume all rivals still competing also have draws 20 points higher (which would imply $q_2 = q_1 + 16$, $q_3 = q_2 + 12$, $q_4 = q_3 + 8$). Nonetheless, this incorporates information about draws x_3 and x_4 to a substantially greater degree than observed in the actual dropout prices p_3 and p_4.

While average responses to the questionnaire are very close to identical between mean and sum conditions, a systematic difference in the distributions underlying these averages may illuminate behavior patterns. The whole distributions of sum condition responses to $q_1(20)$, $q_1(120)$ and $q_4(180)$ each form only a part of the corresponding distribution of mean condition responses. This is shown in Figures 10–12.

The remaining quarter of the mean condition responses form a behaviorally (and perhaps attitudinally) distinct mode. The distinct responses to $q_4(180)$ lie above 145, apparently "aggregating" the information of the draws 120, 140, 160, 180, as if these subjects have no fear to compete to high prices when warranted. The distinct responses to $q_1(20)$ lie below 45, perhaps "hesitating" appropriately when the only information they have is pessimistic. The distinct responses to $q_1(120)$ lie above 125, who "reveal" their supposition that other bidders still competing have higher draws and who are willing to bid high on this basis. Such revealers do not exist in the bidding data, where none of the 32 cases with $x_1 > 120$ have $p_1 > 120$. Notice that when x_1 is low, someone assuming rivals' draws are higher than x_1 cannot be distinguished from someone simply assuming that rivals' draws are typical (i. e., average 100). Similarly, hesitators cannot be separated in cases where $x(1)$ is high, and aggregators cannot be distinguished when x_4 is low.

Behavior which combines "hesitating" and "aggregating" appears significantly more rational than most subjects' behavior observed in the bidding. Responses to questionnaires showing both attitudes occur only for 3 subjects in experiment 2M. The other 6 aggregators are also in the distinct $q_l(120)$ mode, responding as if willing to compete to high prices whether justified or not.

4 Indications

Despite the apparent difficulty subjects faced in making sufficient adjustments for the adverse selection problem in these auctions, we have observed the framing effect of mean versus sum condition nullified in impact on market outcomes. The story across most subjects is one of the framing effect notably altering the first dropout price, but what could be loosely called market forces constraining later dropout prices so as to wipe out this initial impact.

The two formulations for the winner's curse model presented in Section 2 both come closer than the symmetric equilibrium model in predicting the observed average level of profit. The responsive winner's curse model also is consistent with the observed regularity of profits occurring when the asset is more valuable than average, and losses when it is less valuable than average. None of these three models is at all an adequate predictor of observed patterns of individual competition. In particular, all three predict continued impact of private information (the draw) upon dropout prices, when the data clearly reject this hypothesis for the 3rd and 4th bidders to drop out.

The simpler statistical properties of this independent information variant of the canonical common–value auction appear if anything to be less important to the rate of subject learning than the disconcerting role of 100 as a natural anchor. There are enough differences to prevent clear and direct statistical comparisons, but the signs are that subjects in the independent information variant may take even longer to learn to deal with the winner's curse than in the canonical variant (cf. Harstad, Kagel and Levin [1989]).

References

Bikhchandani, Sushil and John Riley, "On Unique Equilibria in Open Auctions," working paper, Department of Economics, University of California, Los Angeles, 1989.

Capen, E. C., R. V. Clapp and W. M. Campbell, "Competitive Bidding in High–Risk Situations," Journal of Petroleum Technology, 1971, v. 23, 641–653.

Cassing, James and Richard W. Douglas, "Implications of the Auction Mechanism in Baseball's Free Agent Draft," Southern Economic Journal, 1980, v. 47, 110–121.

Cox, James C., Bruce Roberson and Vernon L. Smith, "Theory and Behavior of Single–Object Auctions," in Vernon L. Smith (Ed.), Research in Laboratory Economics, v. 2, Greenwich: JAI Press, 1982.

Dessauer, John P., Book Publishing, New York: Bowker, 1981.

Harstad, Ronald M., John H. Kagel and Dan Levin, "Revenue Comparisons in Open– and Sealed–Bid Auctions, with Control for Extraneous Sources of Variation," working paper, Department of Economics, Virginia Commonwealth University, 1989.

Kagel, John H., Ronald M. Harstad and Dan Levin, "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," Econometrica, 1987, v. 55, 1275–1304.

Kagel, John H. and Dan Levin, "The Winner's Curse and Public Information in Common Value Auctions," American Economic Review, 1986, v. 76, 894–920.

Kagel, John H., Dan Levin and Ronald M. Harstad, "Judgment, Evaluation and Information Processing in Second–Price Common Value Auctions," working paper, Department of Economics, University of Pittsburgh, 1988.

Knez, Mark and Vernon L. Smith, "Hypothetical Valuations and Preference Reversals in the Context of Asset Trading," pp. 131–154 in Alvin E. Roth (Ed.), Laboratory Experimentation in Economics: Six Points of View, Cambridge: Cambridge University Press, 1987.

Levin, Dan and Ronald M. Harstad, "Symmetric Bidding in Second–Price, Common–Value Auctions," Economics Letters, 1986, v. 20, 315–319.

Milgrom, Paul R. and Robert J. Weber, "A Theory of Auctions and Competitive Bidding," Econometrica, 1982, v. 50, 1089–1122.

Roll, Richard, "The Hubris Hypothesis of Corporate Takeovers," Journal of Business, 1986, v. 59, 197–216.

Tversky, Amos and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," Science, Jan. 1981, 453–458.

Wilson, Robert, "A Bidding Model of Perfect Competition," Review of Economic Studies, 1977, v. 44, 511–518.

Figure 1: Experiments 1M and 1S
Profit, Observed and Predicted

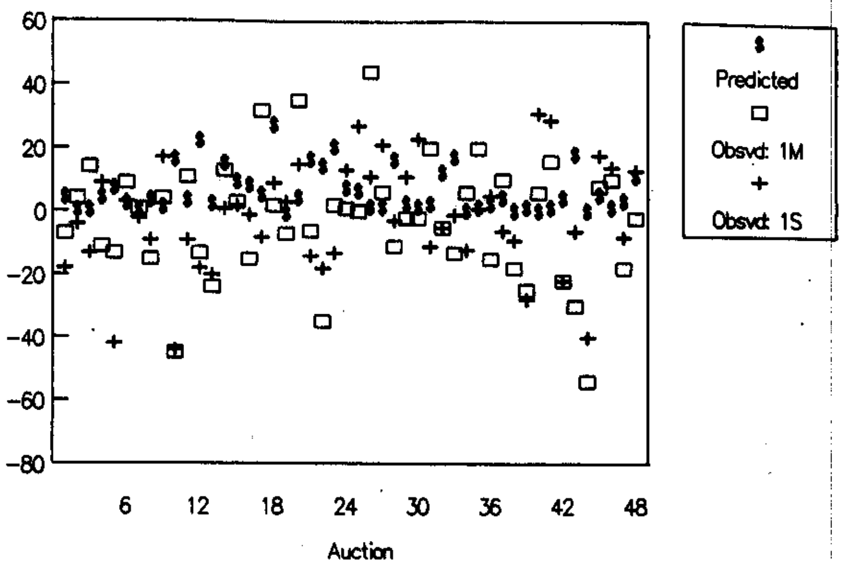

Figure 2: Experiments 2M and 2S
Profit, Observed and Predicted

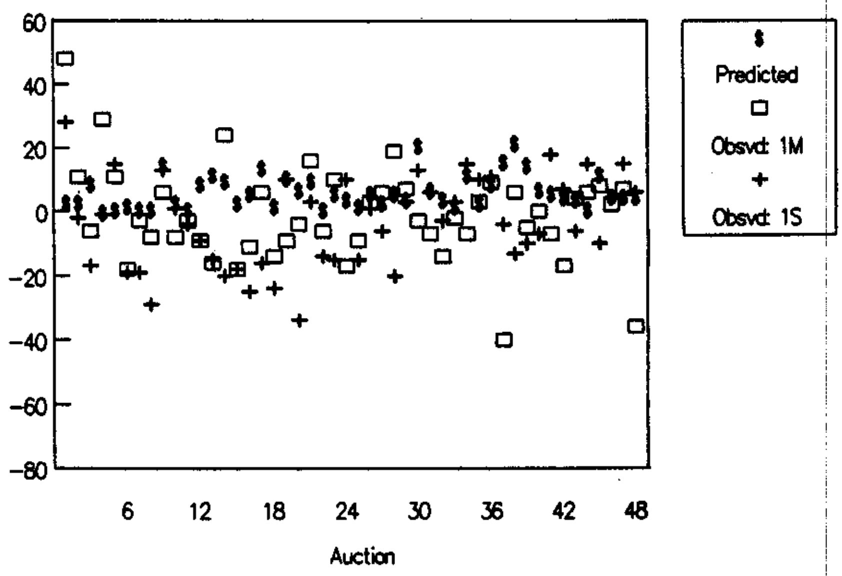

Figure 3: Experiments 3M and 3S Profit, Observed and Predicted

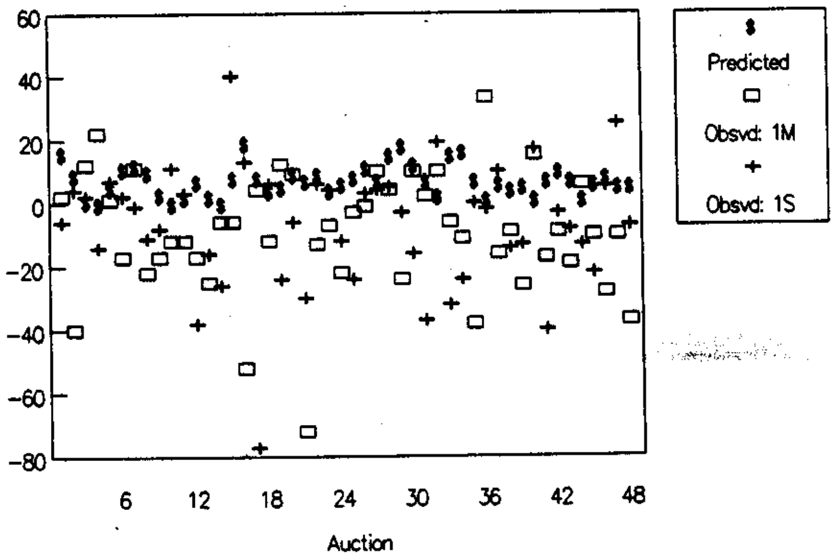

Figure 4: Experiments 4M and 4S Profit, Observed and Predicted

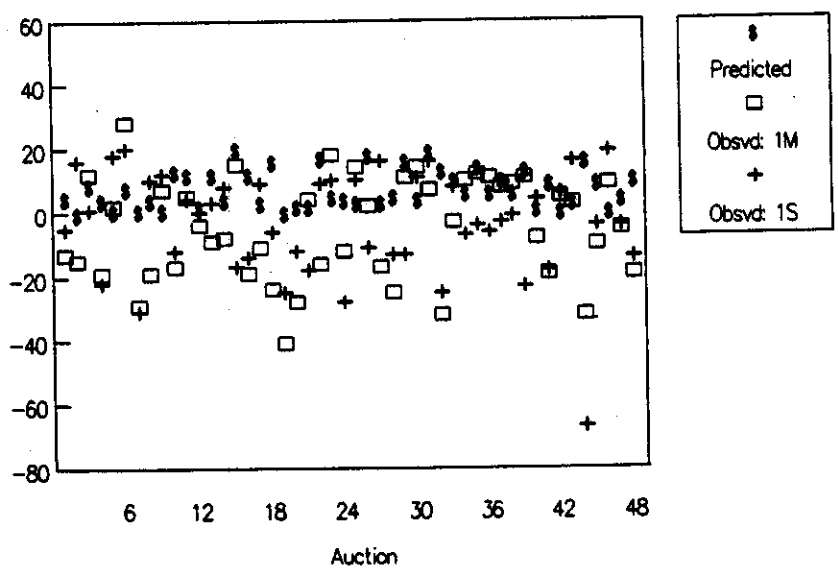

Figure 5: Mean Condition
Profit as a Function of Value

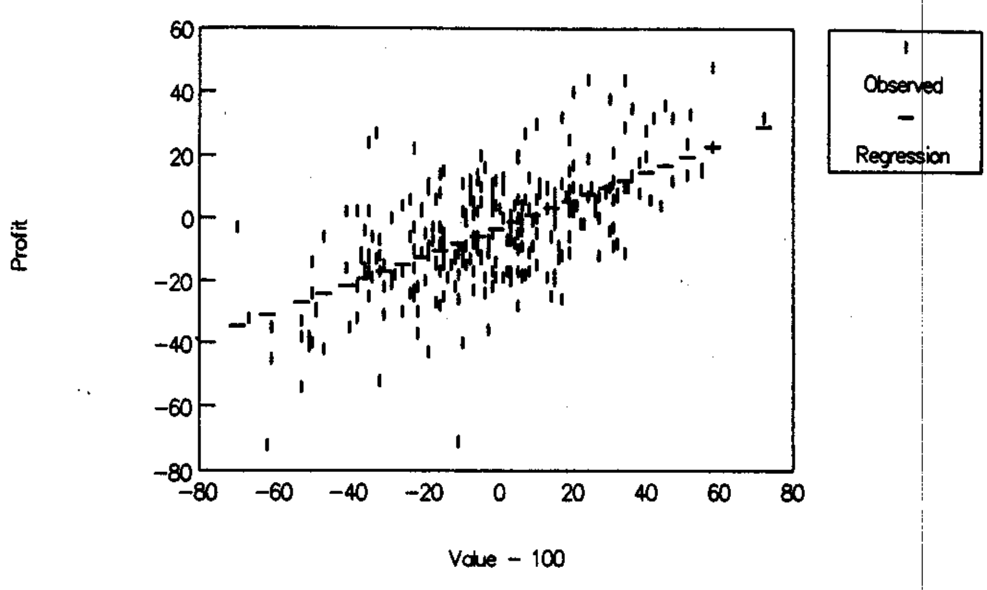

Figure 6: Sum Condition
Profit as a Function of Value

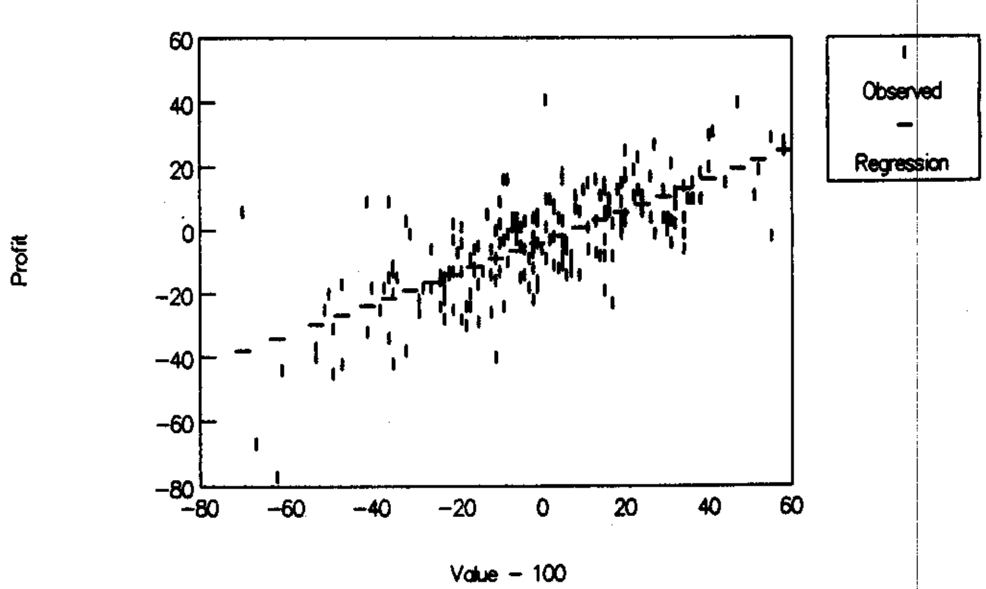

Figure 7: Experiment 2M
Profit as a Function of Value

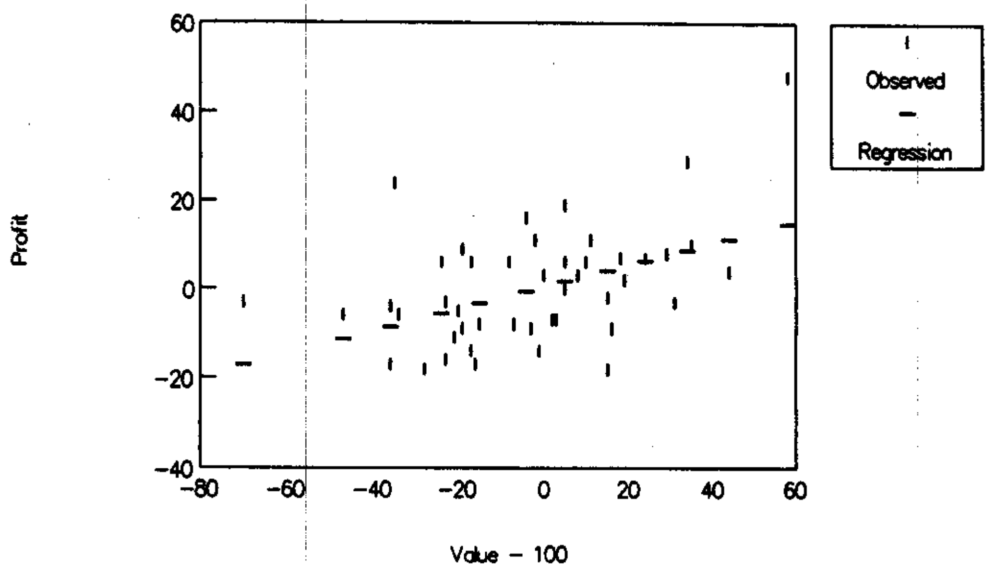

Figure 8: Experiment 3S
Profit as a Function of Value

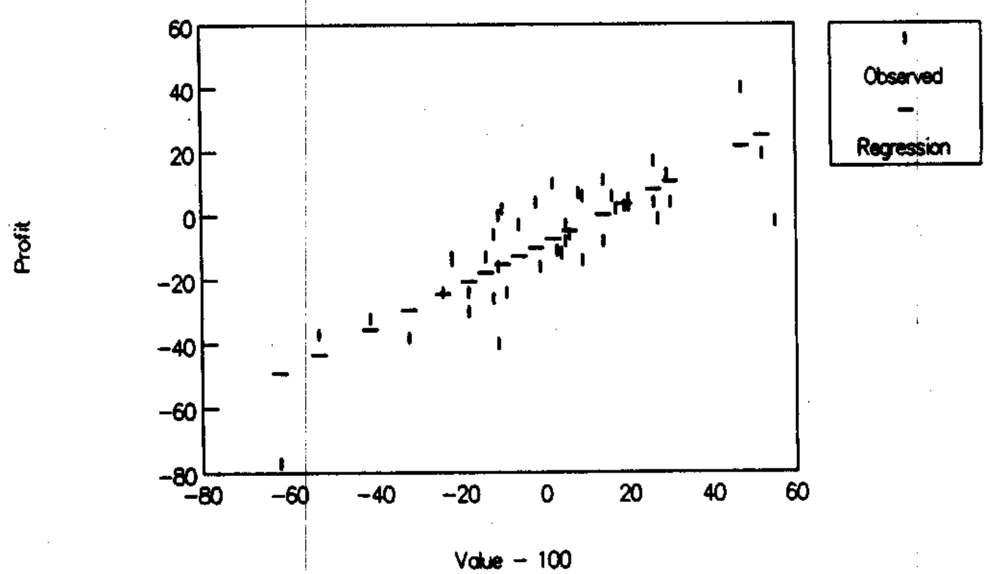

Figure 9: Experiment 4S
Profit as a Function of Value

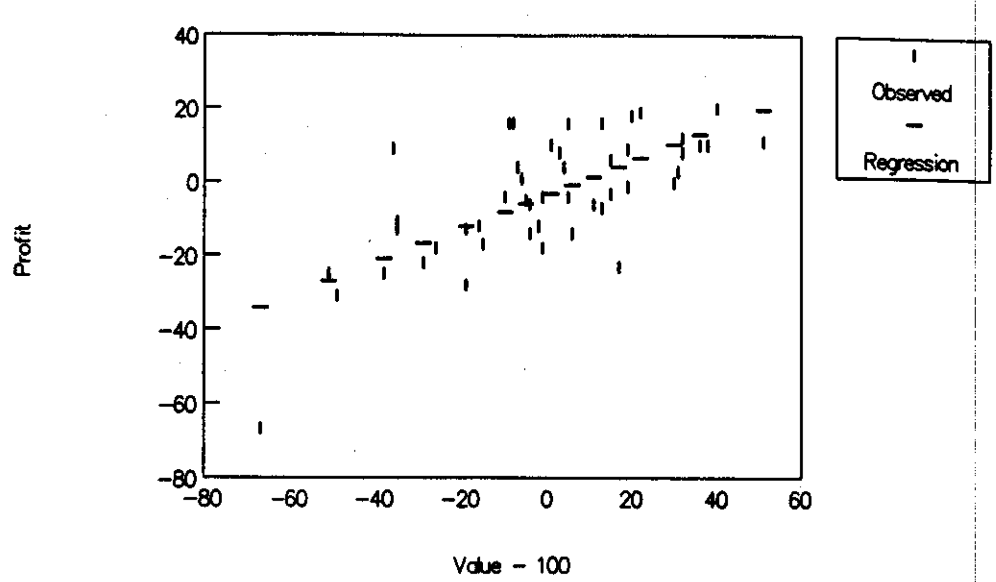

Figure 10: Response Frequencies
1st Drop-Out Price, Draw = 20

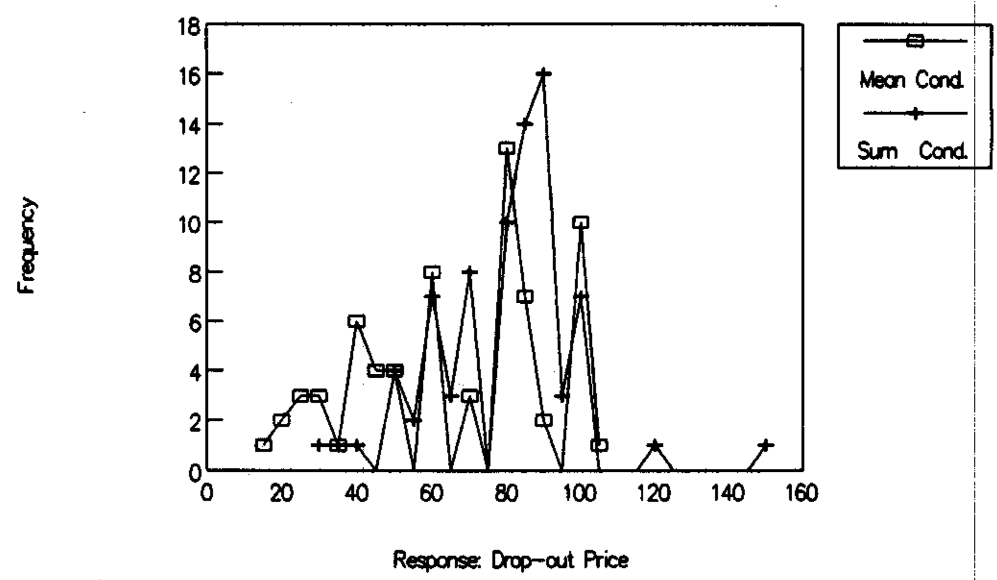

Figure 11: Response Frequencies
1st Drop-Out Price, Draw = 120

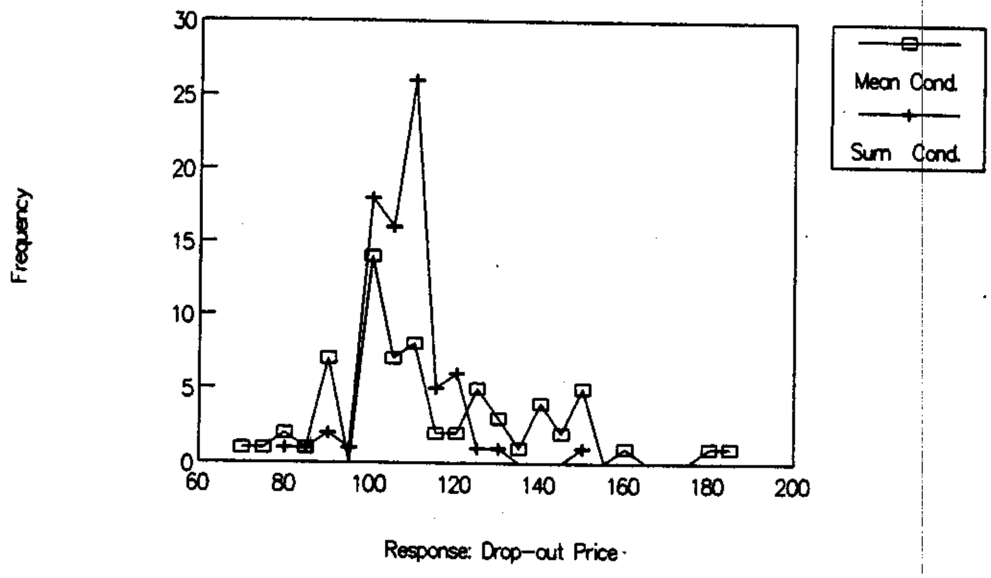

Figure 12: Response Frequencies
4th Drop-Out Price, Draw = 180

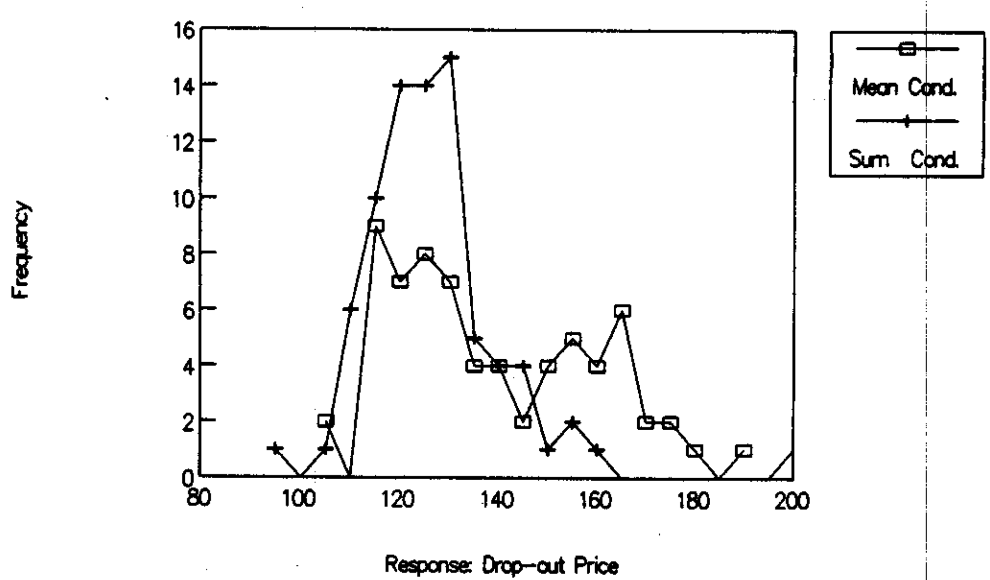

[illegible] Game Equilibrium [illegible]

[illegible]