

Discussion Paper No. 819

**BEHAVIORAL EFFICIENCY II:
A SIMPLE LABORATORY DEMONSTRATION**

Ronald M. Harstad

October 2011

The Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

BEHAVIORAL EFFICIENCY II: A SIMPLE LABORATORY DEMONSTRATION ¹

Ronald M. Harstad
University of Missouri
harstad@missouri.edu
ISER, Osaka University

October 9, 2011

Please consult before citing or circulating

Abstract

Laboratory experiments reporting on shortfalls from allocative efficiency of allocation mechanisms depend on the induced-values methodology, which cannot be extended to the field. Harstad [2011] proposes to observe efficiency of allocation mechanisms without knowing motivations via behavior in appropriately designed aftermarkets. This paper demonstrates the approach in a highly simplified economy: allocation of a single unit of an abstract commodity. In the context studied, second-price auctions are observed to yield significantly greater behavioral inefficiencies than first-price auctions, both in terms of frequency of behaviorally inefficient outcomes, and in terms of the expected size of gains from aftermarket trade missed by the auction itself. The design is shown to be field-ready.

¹ The Economic Science Laboratory at the University of Arizona provided facilities, recruiting and assistance for the conduct of the experiments; thanks to Cathleen Johnson and Lana Sooter for managing this. The J Rhoads Foster Professorship Endowment of the University of Missouri and the Freshwater Group Research Fund of the University of Arizona provided appreciated financial support, and the Institute for Social and Economic Research at Osaka University provided hospitality during the writing. I thank Eric Cardella, Raymond Chiu, Anthony Dubis and Richard Kiser for assistance in programming and conducting the experiments, and Katsunori Yamada for suggestions and assistance in data analysis.

1. Introductory Notes

Recommendations for policy adoption or alteration are more valuable if evidence of the size of shortfalls from allocative efficiency can be provided for the allocation mechanisms or policy instruments under consideration. Such evidence has so far come from laboratory experiments using an induced-values methodology for, e.g., an abstract commodity.² That methodology requires that experimental subjects' motivations are known to the experimenter, and as such is unavailable for field experiments.

In the preceding methodology paper, Harstad [2011], I define *behavioral efficiency*: an outcome of an allocation mechanism is said to be *behaviorally efficient* if an appropriate aftermarket is actually appended to the allocation mechanism and at most a negligible aggregate size of mutually beneficial gains is observed on the aftermarket. That is, a field experiment can first observe behavior under an allocation mechanism (even an informal or culturally-based method of reaching an allocation), and then append a properly constructed aftermarket, and draw inferences from aftermarket behavior as to whether the field experiment's subjects perceived any remaining mutually beneficial transactions. That paper also characterizes construction of an incentive compatible, transparent aftermarket.

This paper is a first demonstration of the concept, providing a concrete example of the appropriate usage of a properly constructed aftermarket to observe allocative efficiencies (or shortfalls therefrom) without relying on knowing subjects' motivations. It finds first-price auctions less behaviorally inefficient than second-price auctions, measures efficiency shortfalls; in this context, subjects' bidding was unaffected by knowing there would be an aftermarket.

2. Laboratory Setting

Five-bidder sealed-bid auctions of a single abstract asset were conducted, in seven sessions (110 subjects) via first-price rules, and in six sessions (85 subjects) via second-price rules.³ There were generally 10, 15, 20, 25 or 30 subjects in the laboratory during a session, with random reassignments into groups of five each period.⁴

Affiliated asset valuations (Milgrom and Weber [1982]) for subjects were determined as follows. In each period, first a random number C , called a *central tendency*, was drawn uniformly from $[\$50, \$1000]$ (all random variables are multiples of $\$0.01$). Then, given a realization c of C , for each subject j an *estimate* X_j was drawn uniformly from $[\$(c - 10), \$(c + 10)]$, conditionally independent. Finally, *asset value* to subject j was $V_j = (3/4)X_j + (1/4)C$; this system incorporates private values (the first term, X_j) to introduce efficiency issues, as well as a natural, small common-value component (C).⁵ These rules were carefully

² A subject j might, for example, be told that she is one of N buyers or M sellers of an abstract good called X that will be traded, with her payoff being the difference between trading prices and induced values, as in "the first unit of X you buy can be resold to the experimenter for $\$8.75$, the second for $\$6.80$, the third for $\$5.10$ " to a potential buyer, or "the first unit of X you sell can be obtained from the experimenter for $\$3.10$, etc."

³ Thus, comparisons across pricing rules are between-subject comparisons.

⁴ Subjects were University of Arizona undergraduates, recruited campuswide via website, and sat at visually isolated computers. A second-price auction session for which less than ten subjects showed was eliminated from data analysis. The experiments were conducted in October and November 2009, using the Z-Tree programming environment (Fischbacher [2007]).

⁵ A principal reason for including a common-value component was to avoid a throw-away bid problem: with independent private values, most values will yield so low a chance of winning as to make serious consideration

explained and examples given. C was not revealed to subjects until end-of-period feedback, which gave a complete, anonymous report of C , X_i 's, V_i 's, all behavior, and profit calculations. The instructions stated that a reserve price, below which the asset would not be sold, was drawn anew before each auction, uniformly from $[(c - 10), (c - 6)]$, and would not be revealed until end-of-period feedback.⁶

Subjects began the experiment with a bank balance of \$12, with profits added and losses subtracted during the session, and the final balance paid in cash. These valuation procedures call for a small winner's curse correction; the 90% confidence interval for the loss in the event a winning bid exceeded the symmetric, risk-neutral equilibrium bid by exactly the winner's curse correction (were all rival bids in equilibrium) is about [\$1.50, \$4.25]. Thus, three to four such losses could likely be handled without the balance going negative.⁷

3. Methodology Implementation

To discern from subjects' behavior whether an auction attained an efficient outcome, the experiment appended an aftermarket designed as follows. Once all subjects had typed and submitted their bids, the winning bidder was determined (throughout by fair random tie breaking if necessary). Then each bidder was informed of the price determined in the auction and whether his bid acquired or did not acquire the asset. Some seconds later, the aftermarket that, before bidding, the subjects had been told would follow the auction was begun.

A price clock ticked up on all subjects' screens, rising by \$0.25 every two seconds (though more slowly in the first period with an aftermarket), beginning at a random price calculated to be acceptable to all subjects but noisy enough to avoid revealing information about the still-unknown C . The bidder who acquired the asset was labeled the *offerer*, and asked to click the "Accept" button on the screen when the price reached the lowest price at which he was willing to sell the asset just acquired in the auction to one of the losing bidders. Each of the four bidders that did not submit the highest bid was asked to do nothing so long as the prices being shown were prices at which he would be willing to buy the asset from the winning bidder, and then to click "Accept" at the highest such price. No subject observed any information about other subjects' behavior in the aftermarket until all five had clicked on a price.

Instructions had carefully described the rules relating these Accept Bids (of the four bidders who did not acquire the asset) and Accept Ask (of the acquirer) to possible aftermarket transactions. [1] If the offerer's Accept Ask exceeds all four Accept Bids, there is no aftermarket transaction. [2] If at least two Accept Bids are no lower than the Accept Ask, the asset is transferred from the offerer to the bidder selecting the highest Accept Bid, at a price set by the second-highest Accept Bid. [3] If the highest Accept Bid exceeds the Accept Ask and it exceeds all other Accept Bids, a random number R , drawn before the

of what to bid not worthwhile. Unfortunately, how high a value a subject has to draw before he or she chooses to pay attention is unobservable. In the current design, all estimates between \$60 and \$990 have the same expected profitability, removing throwaway bid concerns. (Data analysis only includes the cases, almost 99%, where subjects' estimates were in the [\$60, \$990] range.)

⁶ As expected, the reserve price was never binding.

⁷ If a subject's balance became negative, he was given a \$20 loan to be repaid out of his final bank balance. Two of 195 subjects could not quite repay the loan; it was of course forgiven and they were paid only the usual \$5 show-up fee.

auction from multiples of \$0.01 in $[\$c, \$(c + 15)]$ equiprobably, determines the aftermarket outcome. If R falls between the highest Accept Bid and the Accept Ask, the asset is transferred from the offerer to the bidder selecting the highest Accept Bid, at a price set equal to R ; otherwise, there is no aftermarket transaction.

This aftermarket design justifies the inferences about behavioral efficiency to be drawn from observations of aftermarket behavior: any mutually beneficial trade revealed is transacted with positive probability, and in no aftermarket transaction is the price determined by the behavior of either transacting party. Among possible aftermarket designs, prior experimental evidence (Harstad [2000]) suggests the use of the price clock makes aftermarket incentives as transparent as possible. Whenever at least one bidder who did not acquire the asset in the auction selects an Accept Bid above the offerer's Accept Ask, a mutually beneficial trade that the auction did not achieve has been identified (whether or not the aftermarket actually transacts that trade).

4. Session Protocol

Each experimental session ran 150 minutes and followed a multi-phase protocol, to build the desired treatment step-by-step from simpler games. After instructions regarding the whole session and the first phase, that first phase exposed subjects to the software of the aftermarket, without introducing the word "aftermarket." In phase 1 (4-5 periods), each subject was informed of a list of all five private values of the abstract asset (told which was his value), which were drawn i.i.d. uniform on $[\$5, \$10]$. Per instructions, one subject was chosen at random to be the offerer, the others bidders.⁸ As just described, the offerer was asked to click on an Accept Ask, the four bidders to click on Accept Bids. Then the aftermarket rules above were used to determine payoffs for the period, which were simply asset value minus transaction price for the buyer, and transaction price minus asset value for the seller, if there was a transaction, and zero for all non-transacting subjects.

Further instructions were distributed and read before each following phase. Phase 2 (6-7 periods) introduced private information, with subjects' private values first revealed to all group members (anonymously) during end-of-period feedback. Phase 3 (6-7 periods) introduced two changes: [i] all five subjects were now bidders asked to select Accept Bids (that is, in a closed-clock variant of an English auction), and [ii] the private values were now affiliated (as in section 2, except that $V_i = X_i$). Phase 4 (6-7 periods) set aside the price clock, introducing bidding in a sealed-bid auction (first- or second-price, depending on the session). Phase 5 (8-11 periods) introduced affiliated values, the $V_i = (3/4)X_i + (1/4)C$ valuations detailed in section 2.

All this led to the phase of principal interest, phase 6, which re-introduced the software from the first two phases, but with the offerer being the bidder who acquired the asset in the sealed-bid auction, and the following price-clock activity called an aftermarket. Phase 6 was generally limited by the time constraint, 6-11 periods. The session ran faster when there were fewer groups (with the software always waiting for the last subject in the session to bid, to peruse feedback, etc.); in four of the first-price sessions, we were able to run a final phase 7. Phase 7 had aftermarkets only in even-numbered periods, with the sealed-bid auction the

⁸ To generate possible gains from trade frequently, the program chose the subject with the highest, second-highest, ..., lowest private values with probabilities $\{1/8, 1/8, 1/4, 1/4, 1/4\}$.

final determination of period profits in odd-numbered periods. In the other nine sessions, phase 6 was the final phase.

5. Contrasting Predictions

Auction theory predicts aftermarket activity with positive probability following first-price auctions, but with zero probability following second-price auctions. It is straightforward to show that the unique, risk-neutral, symmetric Bayesian equilibrium of auction-cum-aftermarket (either first- or second-price) is to submit one's equilibrium bid in the auction and to select an Accept Bid or Ask in the aftermarket most nearly equal to one's rational Bayesian-updated willingness-to-pay or -accept. In this equilibrium, publicly announcing the price attained in a first-price auction informs each losing bidder (but not the winner) of the amount by which his bid lost. Whenever a bidder lost by a sufficiently small margin, rational updating leads to his willingness-to-pay exceeding the winning bidder's willingness-to-accept (as he knows of a second estimate nearly as high as the winning bidder's estimate, which can be inferred from the price set by the winner's monotonic equilibrium bid function).

No similar occurrence is possible following announcement of the price in second-price auctions. Here the price reveals the private information of the second-highest bidder, who Bayesian updates on the basis of learning that one rival estimate was higher and three lower, and this leads to a willingness-to-pay that exceeds his equilibrium bid, while the winning bidder's updating leads to a willingness-to-accept that is less than his equilibrium bid. However, the second-price auction equilibrium is envy-free: these two adjustments of willingness-to-pay and to accept sum to less than the difference between the two highest bids, and thus do not change their ordinal rank.⁹

To my knowledge, prior auction experiments have either induced private values (independent, as in phases 1 and 2, or affiliated, as in phases 3 and 4) or common values (modifying section 2 so that $V_i = C$, hence there is no inefficiency generated by the bidder with the highest estimate being outbid). Nonetheless, in both settings, bidders have bid significantly above the risk-neutral symmetric Bayesian equilibrium (Kagel [1995]), and (more pertinent here) have exhibited more heterogeneity in this overbidding in second-price than in first-price auctions.¹⁰ Thus, prior laboratory experiment results predict more aftermarket activity following second-price auctions.

6. Aftermarket Observations

First-price [second-price] auctions were observed to be behaviorally efficient in 72% [57%] of the observations (cf. Table 1). In 28% of 203 first-price auctions, and 43% of 142 second-price auctions, at least one bidder who was outbid was observed to be willing to buy the asset from the high bidder for mutual gain. (These percentages naturally sum occurrences where the aftermarket transacted with those where the random price fell below

⁹ Though stated slightly differently, the results in this and the previous paragraph are not new, and can be pieced together from Milgrom [1981], Milgrom and Weber [1982], and Harstad and Bordley [1996].

¹⁰ Kagel and Levin [1986] report on 199 first-price, common-value auctions, and Kagel, Levin and Harstad [1995] on 154 second-price, common-value auctions. To adjust for varying number of bidders, I calculated a statistic for each session that takes the frequency with which the high signal holder was the high bidder and subtracts $1/n$. The weighted (by number of auctions) average of these statistics was 50.93 for first-price auctions and 38.73 for second-price auctions.

the high bidder's Accept Ask or above the one Accept Bid exceeding that Accept Ask.) This difference is significant at the 1% level in a Pearson test.

		Table 1	
		Observations	
		First-Price Auctions	Second-Price Auctions
1	Aftermarkets Observed	203	142
2	Behaviorally Efficient	72%	57%
3	Mean of Shortfalls	\$4.45	\$4.16
4	Shortfall Capacity		\$10.26
5	Aftermarket Fraction		41%
6	Expected Shortfall	\$1.24	\$1.79

Where aftermarket behavior exhibited such gains, the difference between the most an outbid bidder will pay and the least the high bidder will accept is a behavioral measure of the shortfall from efficiency, averaged in row 3. While shortfalls when observed were larger in first-price auctions, when zero shortfalls are averaged in for the behaviorally efficient outcomes, the expected shortfall in row 6 is smaller for first-price auctions.

It bears emphasis that, while these auctions sold induced-value assets, the behavioral efficiency and shortfall measures make no use of any information contained in the induced values. These reports stem solely from subjects' behaviors: their sealed bids, Accept Bids and Accept Asks, and in no way depend on any information about subjects' motivations.

Reports of allocations reached in induced-values experiments can provide efficiency measures in percentages, because the dollar value of total gains from trade in Pareto-efficient allocations can be calculated from the induced values. This methodology cannot be used in the field.

In some situations, behavior in the original allocation mechanism can offer a benchmark for the economic significance of the size of shortfalls from efficiency. This experiment demonstrates both the possibility and its limitations.

Second-price auctions are incentive compatible, in that the bid selected determines only whether the bidder wins or not; the price is solely determined by the highest rival bid. In particular, the risk-neutral symmetric equilibrium bid is the expected asset value conditioned on an assumption that the bid is pivotal.¹¹ For the distributions of section 2, this implies that bids should differ from bidders' expected values by a constant.¹² Hence, differences between two bids submitted in second-price auctions should be equal to the differences between the two bidders' willingnesses-to-pay and therefore measure the gain from trade if

¹¹ For private values, this is the dominant strategy discovered by Vickrey [1961]. This feature of second-price, common-value auctions was first found by Matthews [1977]; the intuition is presented in Harstad and Bordley [1996].

¹² This neglects estimates in the ranges [\$50, \$60] and [\$990, \$1000], for which the difference is not constant.

the asset were hypothetically to be transferred from the lower bidder to the higher bidder. These observations include 586 differences between a losing second-price bid and the high bid in the same auction. The average of those 586 bid differences is the \$10.26 reported in row 4 above. This is the *capacity* for shortfalls from efficiency in the sense that, had the original auctions allocated the asset equiprobably among inefficient acquirers, aftermarkets that reallocated to the efficient acquirer could average \$10.26 in gains from trade unattained by such a complete misallocation to a random inefficient acquirer.

The best measure I can envision to attach an economic significance to the \$4.16 mean of shortfalls revealed by aftermarkets following second-price auctions is that it is 41% of the shortfall capacity.

The absence in Table 1 of a comparable benchmark for first-price auctions is not an oversight. Differences between a losing first-price bid and the high bid in the same auction could be averaged. However, without the incentive compatibility of second-price auctions, these first-price bid differences have no similarly strong argument to measure gains from hypothetical transfers between the bidders: optimizing the {profitability given winning/probability of winning} tradeoff in the risk-neutral symmetric equilibrium of the first-price auction yields a nonlinear term in the bid function (corresponding to the term that is a constant in second-price auctions). In that equilibrium, if bidder A outbids bidder B by \$8, A 's willingness-to-pay exceeds B 's, but the \$8 bid difference is not a measure of the willingness-to-pay difference.

7. Submitted-Bid Impact

The unique risk-neutral symmetric Bayesian equilibrium of the game consisting of the sealed-bid auction (either pricing rule) followed by the aftermarket is for each bidder to make the same equilibrium bid as if there were no aftermarket, and then truthfully reveal in the aftermarket. Despite the theory, it is an empirical question whether subjects bid the same way when they know an aftermarket will follow; there might be reasons subjects would find for bidding less, or for bidding more, in an auction when knowing there will be an aftermarket.¹³ The protocol in section 4 is designed to shed light on this question.

The following linear bid function was estimated separately from the first-price and second-price data:

$$M_{st} = \text{const} + \beta_x \text{Exper}_{st} + \beta_a \text{After}_t + \text{error}_{st},$$

where the markup M_{st} was the observed bid minus the asset value estimate X_{st} for subject s in period t ; Exper_{st} was a control for possible learning effects, the number of periods of experience in the affiliated-values auctions; After_t was a dummy variable taking the value 1 if the subject knew the auction in period t would be followed with an aftermarket, 0 if the subject knew the auction would not be followed with an aftermarket.

Estimates obtained from OLS linear regressions with clustering by subject are presented in Table 2. For both types of auction rules, a null hypothesis that subjects bid no differently

¹³ Among the possibilities are that a subject might perceive an opportunity to win the auction profitably and then profit further by selling in the aftermarket, which could be perceived as suggesting more aggressive bidding than if there were to be no aftermarket; or that a subject might perceive the aftermarket as a second chance to obtain the asset, which could be perceived as suggesting less aggressive bidding than if there were to be no aftermarket.

when knowing there would be an aftermarket as when there would not cannot be rejected at anything vaguely approaching conventional levels of significance.¹⁴

Table 2		
Estimates		
	First-Price Data	Second-Price Data
const	-2.346	-1.435
Std. error	(0.567)	(1.629)
Significance	0.001	0.38
β_x	0.045	0.393
Std. error	(0.075)	(0.370)
Significance	0.55	0.29
β_a	-1.242	-5.063
Std. error	(1.431)	(4.738)
Significance	0.39	0.29
# Observations	2215	1075
F test:	1.24	0.58
Significance	0.293	0.562

8. Distinct Measure?

What can be said about how well behavioral efficiency tracks Pareto efficiency? As these experiments used induced values, they can yield insights into the differences between these measures. That is, assume (critically) subjects are all risk-neutral (or identically risk averse), and assume completeness of the induced motivations (in particular, assume away interdependent preferences, nonpecuniary preferences, and satiation in cash). Then in a Pareto-efficient allocation, the asset is acquired by the subject with the highest estimate.

In most observations, when either a first-price or a second-price auction reached a Pareto-efficient allocation, behavioral efficiency was observed in the aftermarket, and the inverse: Pareto-inefficient auctions led to behaviorally inefficient outcomes, mutual gains observed in the aftermarkets.

The two distinctions from tracking were both observed in significant minorities of the observations. [i] In 15% of first-price auctions and 24% of second-price auctions, the

¹⁴ Failure to reject this null was found in alternatives that did not cluster or added subject fixed effects; alternative where the bid was the dependent variable and asset value estimate an independent variable were nearly identical.

efficient acquirer was the high bidder, so the outcome is assumed Pareto-efficient, yet the aftermarket found a mutual gain could arise from a transfer to an outbid rival with a lower estimate of asset value. [ii] For both first-price and second-price auctions, 16% of observations found an inefficient acquirer submitting the high bid and then clicking on an Accept Ask that exceeded all Accept Asks, including that selected by the efficient acquirer whose estimate exceeded his.

While some variant of an endowment effect could lead to the second way in which behavioral efficiency has been found distinct from Pareto efficiency, it bears notice that distinction [i] is completely inconsistent with an endowment effect. More importantly, being able to observe which auction outcomes are Pareto efficient and thus observe these distinctions depends on having induced values and assumed motivational completeness and identical risk tolerances. Using aftermarkets to observe the size and frequency of shortfalls from behavioral efficiency requires none of these.

9. Readiness Remarks

These experiments used carefully designed aftermarkets to observe behavior yielding the inference of shortfalls from efficiency in 203 first-price auctions with a 28% frequency and \$1.24 expectation, and in 142 second-price auctions with a 43% frequency and \$1.79 expectation. Adding the aftermarket did not in these experiments affect bidding in the original auction. While the experiments employed induced values, no aspect of induced values was utilized in reaching these conclusions. A field study of a single-asset auction could exactly mimic these procedures to obtain evidence on whether the efficient acquirer won the auction and if not, the size of the inefficiency that arose, even in cases where existence of equilibrium is in doubt (cf. Jackson [2009]) or equilibrium is incalculable.

The random-price procedure used when exactly one outbid bidder selected an Accept Bid above the Accept Ask could directly be used to observe shortfalls from efficiency in bilateral bargaining over transfer of an indivisible asset or service. As pointed out in Harstad [2011], the budget-balance feature of the aftermarkets reported here (that the amount the aftermarket purchaser paid is exactly the amount the aftermarket seller received) is not essential to appropriate aftermarket design. When field experiment budgets allow, experimenters can provide incentives for field subjects to participate in aftermarkets by designing aftermarkets that run an experimenter-covered deficit. For example, in bilateral bargaining, efficiency conclusions could be drawn from an aftermarket design that, whenever announcements indicated a mutual gain, [a] paid the aftermarket seller 5% more than the random price and charged the aftermarket buyer 5% less than the random price, or [b] paid the aftermarket seller the Accept Bid and charged the aftermarket buyer the Accept Ask.

Fairly straightforward complications of the aftermarket design used here can accommodate observing efficiency shortfalls for mechanisms seeking to allocate multiple homogeneous assets. For example, following a mechanism for allocating two homogeneous assets, potential buyers in an aftermarket could be asked for a pair of Accept Bids if seeking to buy, an Accept ask if one asset won, or a pair of Accept Asks if two, with the rules that whenever an Accept Bid by a rival fell between an Accept Bid and a lower Accept Ask, it set the price for that transaction, and a random price was consulted when necessary. It would not matter whether bidders had single-unit or multi-unit demands. In larger, semi-competitive markets for homogeneous assets, a variant on a call market could serve as an aftermarket (so long as no trader were seeking both to buy more and to sell some of what he

had obtained), with the highest quantity where the demand price exceeded the supply price transacted, buyers paying the price of the last accepted supply unit, sellers receiving the (higher) price of the last accepted demand unit, and the experimenter covering the deficit.

In principle, aftermarkets could yield behavioral observations of shortfalls from efficiency even in cases where public goods were being allocated, externalities arose, services had time-dependent valuations, and/or packages of goods had synergistic values. It would still be the case that a proper aftermarket design would require that the focal equilibrium of the allocation-mechanism-cum-aftermarket be focal equilibrium behavior in the allocation mechanism followed by truthful revelation in the aftermarket. Exactly what constraints these issues pose on aftermarket design is beyond the scope of this first demonstration.

References

- Urs Fischbacher [2007], z-Tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* 10, 171–78.
- Ronald M. Harstad [2000], Dominant strategy adoption and bidders' experience with pricing rules, *Experimental Economics* 3, 261–80.
- Ronald M. Harstad [2011], Behavioral efficiency I: a definition and methodology, Economics working paper, University of Missouri.
- Ronald M. Harstad and Robert F. Bordley [1996], Lottery qualification auctions, pp. 157–83 in: M. R. Baye (ed.), *Advances in Applied Micro-Economics*, Volume 6, *Auctions*, JAI Press, Greenwich, Ct.
- Matthew O. Jackson [2009], Non-existence of equilibrium in Vickrey, second-price, and English auctions, *Review of Economic Design* 13, 137–45.
- John H. Kagel [1995], Auctions: a survey of experimental research, pp. 501–86, in John H. Kagel and Alvin E. Roth, (eds.), *The Handbook of Experimental Economics*, Princeton University Press.
- John H. Kagel and Dan Levin [1986], The winner's curse and public information in common-value auctions, *American Economic Review* 76: 894–920.
- John H. Kagel, Dan Levin and Ronald M. Harstad [1995], Comparative static effects of number of bidders and public information on behavior in second-price common-value auctions, *International Journal of Game Theory* 24, 297–319.
- Steven Matthews [1977], Information acquisition in competitive bidding processes, working paper, Humanities and Social Science, California Institute of Technology.
- Paul R. Milgrom [1981], Rational expectations, information acquisition, and competitive bidding, *Econometrica* 49, 921–43.
- Paul R. Milgrom and Robert J. Weber [1982], A theory of auctions and competitive bidding, *Econometrica* 50, 1089–1122.

