

Implications of Approximate Equilibrium Concepts in Sealed-Bid Auctions¹

Theodore L. Turocy III²

Department of Managerial Economics and Decision Sciences

J. L. Kellogg Graduate School of Management

Northwestern University

2001 Sheridan Road

Evanston IL 60208

arbiter@kellogg.nwu.edu

February 28, 2001

¹I would like to thank Mark Satterthwaite, Roger Myerson, Bob Weber, Jim Dana, Richard McKelvey, Dan Levin, and John Kagel for helpful comments and suggestions, as well as feedback from seminar participants at Duke, IBM Watson Labs, Texas-Austin, Texas A&M, SUNY-Albany, Ohio State, and Miami. All errors are, of course, my own.

²After August, 1, 2001, Department of Economics, Texas A&M University, College Station TX

Abstract

This paper analyzes the structure of the expected utility functions of bidders in sealed-bid auctions, and considers how payment rules interact with the effects of small optimization errors by bidders. The range of expected revenues to the seller from ε -equilibrium profiles expands rapidly for small values of ε , indicating that the seller revenue prediction of Bayes-Nash equilibrium may not be robust in the presence of optimization errors by bidders. Numerical estimates show an error on the order of 0.5 percent in expected value terms by bidders can yield a 10 percent change in the amount of seller revenue. Furthermore, the range of seller revenues supported by the set of ε -equilibria is smaller in the first-price auction than in its second-price counterpart. Computation also indicates that the second-price auction has ε -equilibria, for given ε , which are more inefficient than the ε -equilibria of the first-price auction. These results are related to observations from the experimental literature on auctions.

1 Introduction

The design and analysis of auctions has been of interest in various fields of economics for several decades now, dating back to the seminal contributions of Vickrey [20], and Milgrom and Weber [16], among others. Recently, the design of auctions and related mechanisms has gained practical relevance in applications as varied as the allocation of spectrum licenses and the procurement process of companies in business-to-business electronic exchanges.

Economists and game theorists design such institutions with specific objectives in mind. In some cases, efficiency is desired; in others, revenue maximization for the seller; in yet others, specific distributional policy goals may come into play, such as encouraging bidding by minority-owned firms. For a given institution, the goal is to make predictions about the behavior of the players (whether they be individuals, firms, or other agents) so as to be able to predict the outcome. The traditional approach to this is to invoke a suitable version of Nash equilibrium as the solution concept for the game.

Underlying this all, however, is the knowledge that the agents who actually play these games have limits on their rationality. They may imperfectly understand the rules of the game; they may have inconsistent beliefs about the private information of others; they may not be able to accurately compute their best strategic choices, given their information. Any of these, in practice, will lead to non-Nash play, and most likely non-Nash outcomes.

In view of this, when comparing the properties of possible mechanisms in a particular setting, I argue that it is necessary to not only compare the predictions of Nash equilibrium itself, but also to make an analysis of the robustness of these predictions to deviations from equilibrium play. In this paper, I will approach this robustness question from the perspective of modeling agents who are not necessarily able to choose their optimal play due to computational limitations. Instead, the agents will be assumed to choose actions which are “good enough”, by some criterion.

In his comparison of auction institutions, Vickrey extolled the praises of the second-price mechanism as being “advantageous to all parties”. He wrote,

In cases in which, by reason of asymmetry among the bidders, errors in evaluation, or mistakes in strategy, the result with the “top-price” method is non-optimal, a change to the “second-price” method will yield an increase in the aggregate profits to be shared among seller and buyers.¹

Vickrey continues in the paragraph to talk through some of the efficiency implications of asymmetry among the bidders. Unfortunately, he does not go into further detail about what he means by “mistakes in strategy”, but in the following paragraph he posits a benefit of the “simplicity” of the dominant-strategy equilibrium of the second-price mechanism:

¹Vickrey [20], p. 21.

In addition to the gain from the improved allocation of resources, there is another possible gain that is not covered in the above analysis, which abstracts from the costs involved in the negotiations. In the top-price method of negotiation, as in the Dutch auction, bidders, in order to maximize their expectation of profit, must concern themselves not only with their own appraisal of the article but also with their estimate of the value that others will place on it and their expectation of the bidding strategy that others will follow. ... It is one of the salient advantages of the second-price method that it makes any such general market appraisal entirely superfluous... Each bidder can confine his efforts and attention to an appraisal of the value the article would have in his own hands, at a considerable saving in mental strain and possibly in out-of-pocket expense.²

I interpret this paragraph as discussing not only the dominant-strategy nature of the equilibrium, but also the ease in which bidders can map from their private value into their optimal bid. In this conception, bidders find the optimization problem to be very simple, and as such, the second-price mechanism should be more robust.

This paper tests the robustness of two well-known auction environments under the assumption that players may make errors which have “small” negative payoff effects. Even though each player may be making only a small optimization error, the interaction of the errors may have substantial effects on the outcome; furthermore, the magnitude of these effects may depend upon the details of the mechanism under study.

The subject of the robustness of equilibrium to bidder error has been broached in a paper by Bajari [2], in which he applies the quantal response equilibrium concept of McKelvey and Palfrey [14] to bidding behavior. The quantal response equilibrium describes behavior in an environment in which bidders make errors in computing their payoffs. Bajari focuses on the issues raised for econometric analysis of bidding data; in contrast, this paper addresses the effects of optimization error upon revenue and efficiency predictions.

The approximate equilibrium concept that will be applied in this paper is ε -equilibrium, as introduced by Radner [17]. In an ε -equilibrium, each player plays a strategy which gives him a payoff within ε of the payoff yielded by his best reply to his opponents’ strategies. Radner’s definition of his concept in terms of the “payoff space”, rather than the “message space”, to use mechanism design terminology, emphasizes that the objective of the agents is not to play according to a Nash profile, but rather to do as well as possible for themselves, given what their opponents are doing.

Radner proposed the following interpretation of ε -equilibrium:

One type of answer [to the question of how to interpret ε] refers to the various

²Vickrey [20], pp. 21-22.

costs of discovering and using alternative strategies, and alludes to the possibility that a truly optimal response might be more costly to discover and use than some alternative, “nearly optimal” strategy. In this interpretation, the “epsilon” for a particular firm represents a judgement of the firm that the additional benefits from improving its strategy would be outweighed by the additional costs. ... It would be consistent with the spirit of the model for this judgement to be in part subjective, rather than necessarily based on some precise calculation of benefits and costs.³

Turning to the experimental literature, an empirical regularity in first-price independent private values auctions is that subjects in the laboratory frequently overbid relative to the risk-neutral Nash equilibrium prediction. Several theories have been advanced to account for this: it is consistent, for example, with risk-averse bidders (even though this phenomenon seems to persist when risk neutrality is induced in the subjects), and is also consistent with a utility benefit for winning the auction.

In a somewhat controversial article, Harrison [9] critiqued the Cox, et al (see, for example, [4]) series of auction experiments by arguing that the deviations from risk-neutral Nash equilibrium bidding could be explained by looking at the payoff space, rather than the message space. In dollar terms, even relatively large deviations from risk-neutral Nash behavior result in only a small opportunity cost to the bidder, relative to bidding according to the best response; thus, this is referred to as the “flat-maximum” critique. Harrison proposes a metric he names the “Nash conceptual experiment”, in which he evaluates the strength of a bidder’s incentive to adhere to Nash play if the bidder assumes his opponents all are bidding according to the Nash profile.

Harrison’s article led to a spirited exchange in the *American Economic Review* a few years later ([13], [6], [15], [10]). It is uncertain to what degree one can practically apply Harrison’s argument to the experimental evidence, since the “flatness” of the maximum can be altered by affine transformations of the agents’ utility functions. However, Harrison’s idea can be applied to compare the implications of the behavior of the same agents across different mechanisms, since one can speak of one mechanism having a “flatter” maximum than another. It is this comparative approach that this paper adopts.

During the *AER* exchange, Merlo and Schotter [15] examined experimental data from agents behavior when faced with a decision problem, and characterized participants as “theorists” or “experimentalists”, based upon whether they tended to pick what they perceived to be an optimal point and stick with it (the “theorists”), or whether they adjusted their choices and attempted to learn the shape of their payoff function (the “experimentalists”).

³Radner [17], p. 153.

It is this latter type of behavior which ε -equilibrium attempts, in one particular way, to capture and formalize. An interesting feature of this idea is that whether a particular agent is a “theorist” or “experimentalist” in a particular environment may be endogenous, in the sense that some agents might find some problems to be directly solvable, and adopt the theorist approach, whereas those same agents might find others more difficult, and adopt the experimentalist approach to search for their optimal play.

The remainder of the paper is organized as follows. Section 2 lays out the auction model, and the definition of the ε -equilibrium concept. Section 3 states the optimization problems to be solved, and provides some characterizations of the solutions to these problems. Section 4 describes the algorithms used, and reports the results of numerical analysis of the optimization problems for the case of two bidders. Section 5 discusses the implications of the results for both the design of experiments, and for the development of approximate equilibrium concepts. Section 6 concludes.

2 Definitions and Model

2.1 Auctions

In what follows, the focus will be on symmetric, sealed-bid auctions of the form studied by, for example, Milgrom and Weber [16]. In these auctions, one agent (the seller) wishes to sell one unit of an indivisible good, for which he has zero value. There are two bidders, indexed by $i = 1, \dots, n$, each of whom simultaneously submits a single, nonnegative bid for the object. These bidders are assumed to have private (idiosyncratic) valuations for the good which are distributed uniformly upon the unit interval $[0, 1]$, and realizations of these values are independent across the bidders. The valuation x_i of bidder i is private information.

The seller sells the object to the highest bidder. In the case of a first-price auction, the highest bidder pays his own bid; in the second-price auction, the highest bidder pays the amount of the second highest bid. Ties are broken randomly, although under the assumptions that will be made, they occur with probability zero. The bidder i who wins the object and pays p gets utility $x_i - p$; the utility of losing is normalized to zero.

The set strategies for each bidder i , denoted Σ_i , is taken to be the set of continuous, strictly-increasing functions from the set of types $[0, 1]$ to the set of possible bids. The notation $b_i(\cdot)$ represents one such bidding function for bidder i , the inverse of which will be denoted $\lambda_i(\cdot)$. The expected utility to bidder i to following the strategy profile b_i , when his opponents follow b_{-i} , is written $U_i(b_i, b_{-i})$. (When speaking in terms of inverses, the expected utility will be written $U_i(\lambda_i, \lambda_{-i})$.) The unsubscripted symbol Σ will refer to the Cartesian product of the Σ_i .

In the unique symmetric Nash equilibrium under the first-price rule, each bidder bids as a function of his valuation x (suppressing subscripts due to the symmetry)

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

that is, he reduces his bid lower than his value to trade off the probability of winning against the price he pays when he wins. In the second-price auction, the bidder never affects price, and thus there exists an equilibrium in which he truthfully reveals his type,

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

Furthermore, in this private-value setting, as Vickrey pointed out, it is a *dominant strategy* to bid one's own type in the second-price auction.⁴

In this independent private values environment where the seller has a value of zero for the object, the expected gains from exchange are $\frac{n}{n+1}$, which is the expected value of the highest order statistic of the n private values of the bidders. The symmetric equilibria given in (1) and (2) result in the same distribution of the gains from exchange. In each case, the expected revenue to the seller is $\frac{n-1}{n+1}$, with the remaining $\frac{1}{n+1}$ split evenly among the n buyers, who each receive $\frac{1}{n(n+1)}$. These equilibria are efficient, with the full expected gains from exchange being extracted.

The motivation to restrict strategies to being strictly increasing is twofold. In general, experimental data suggest that bidders in independent private values auctions do follow increasing bid functions, even though they may not bid according to the risk-neutral Nash equilibrium (see, for example, the data in Cox, et al [5], as reproduced in Davis and Holt [7]). From a more practical point of view, the ability to cast the computations entirely in terms of inverse bid functions permits many of the integrals involved to be evaluated precisely, rather than numerically.⁵ This choice will also prove convenient in characterizing the nature of the ε -equilibria which maximize seller revenue, since the expected values and revenues will vary continuously under small perturbations of the inverse bid functions.

2.2 Definition of ε -equilibrium

The approximate solution concept used in the subsequent analysis is Radner's [17] concept of ε -equilibrium. This notion operationalizes the failure of agents to fully optimize in the following definition.

⁴There exists a continuum of other Nash equilibria in the second-price auction. This equilibrium will be employed as a baseline for the calculations in keeping with the spirit of Vickrey's analysis.

⁵See Appendix B for details.

Definition 1 (ε -equilibrium) *A strategy profile $(\sigma_i)_{i=1}^n$ is an ε -equilibrium, for $\varepsilon \geq 0$, of an n -player strategic-form game, if*

$$U_i(\sigma', \sigma_{-i}) - U_i(\sigma_i, \sigma_{-i}) \leq \varepsilon, \forall \sigma' \in \Sigma_i, \forall i \in \{1, \dots, n\}. \quad (3)$$

That is, the profile σ is such that no player i could expect to gain more than ε by deviating from σ_i to any of his other feasible strategies.

In applying this definition to the auction model, I will take the ex-ante perspective: ε is given in terms of the loss in expected value of following a bidding strategy b_i , relative to the best-response strategy, prior to learning one's own valuation x_i . An alternative definition of ε -equilibrium would take an interim perspective, in which the utility loss incurred by each type x_i of bidder i , is constrained to be no more than ε . Because the motivation of this analysis is to look at the problem from the perspective of a mechanism designer, the ex-ante perspective seems to be preferable. The particular shapes of the computed profiles exhibited later do depend on this choice of definition.

3 Theory

Explicit characterization of the whole set of ε -equilibria would be a difficult task. Instead, I restrict my attention to finding members of the set which result in the most extreme deviations from the Nash predictions, as measured by quantities of economic interest. This paper focuses on the expected revenue to the seller, and the efficiency of the allocation, as measured by the ex-ante percent of gains from exchange realized. Since symmetry and efficiency are closely related in the Nash equilibria of these auctions, efficiency measures how asymmetric ε -equilibria can be. Proposition 6 below shows the relationship between symmetry and maximal seller revenues. Thus, taken together, these measures speak to the extent of the set of ε -equilibria.

In each case, there exists a range of values supported by the set of ε -equilibria. I will avoid speaking of the "distributions" of these values over the ranges, since doing so would imply some distribution over the set of ε -equilibria, and therefore some distribution over the optimization errors being made. Since there exists a continuum of ε -equilibria for positive ε , there is a coordination problem among bidders to choose a particular ε -equilibrium from the set. The details of a particular model of this coordination process would affect the likelihoods of some ε -equilibria being chosen over others. While extensions in the direction of modeling this process are attractive, this paper restricts attention to computing the feasible ranges, which would be the (maximal) support of the distributions over these values. These ranges are interesting in and of themselves; for example, a sufficiently uncertainty-averse seller would

prefer a smaller interval of expected revenues (with the same expectation) over one with a larger support.

3.1 Effects on seller's expected revenue

The principle of revenue equivalence states that the expected amount of revenue the seller anticipates from either a first- or second-price rule in the independent private values setting will be the same; namely, a revenue of $\frac{n-1}{n+1}$ in either case. The prediction for revenue, given a set of ε -equilibria for positive ε , is a range of revenues around the Nash equilibrium prediction. The goal is to find the endpoints of this range. Writing $R(\lambda)$ to be the expected revenue to the seller when bidders follow the strategy profile λ , the optimization problems to locate the endpoints are expressed as follows.

Problem 2 (Maximizing seller revenue)

$$\begin{aligned} \bar{R}(\varepsilon) &= \max_{\lambda \in \Sigma} R(\lambda) \\ \text{subject to } & U_i(\lambda_i^*(\lambda_{-i}), \lambda_{-i}) - U_i(\lambda_i, \lambda_{-i}) \leq \varepsilon \quad \forall i \in \{1, \dots, n\} \end{aligned} \quad (4)$$

Problem 3 (Minimizing seller revenue)

$$\begin{aligned} \underline{R}(\varepsilon) &= \min_{\lambda \in \Sigma} R(\lambda) \\ \text{subject to } & U_i(\lambda_i^*(\lambda_{-i}), \lambda_{-i}) - U_i(\lambda_i, \lambda_{-i}) \leq \varepsilon \quad \forall i \in \{1, \dots, n\} \end{aligned} \quad (5)$$

I now show that any profile λ which is a solution to the problems in displays (4) and (5) must necessarily be symmetric, that is, $\lambda_i = \lambda_j$ for all i, j . To show this, I introduce a relaxed version of these two problems, in which the n constraints imposed by the restriction that each bidder must be making an error in expected value of less than ε is replaced by the one constraint that the *sum* of these errors across bidders is less than $n\varepsilon$. This relaxed version is stated for the maximization case in (6); the minimization version is analogous.

Problem 4 (Maximizing seller revenue, relaxed version)

$$\begin{aligned} \bar{R}(\varepsilon) &= \max_{\lambda \in \Sigma} R(\lambda) \\ \text{subject to } & \sum_{i=1}^n U_i(\lambda_i^*(\lambda_{-i}), \lambda_{-i}) - U_i(\lambda_i, \lambda_{-i}) \leq n\varepsilon \end{aligned} \quad (6)$$

The set of profiles defined in (6) is a superset (a strict superset for $\varepsilon > 0$) of that defined in (4). So, if there exist solutions to the problem in (6) which are feasible in the problem in (4), then those solutions must also be optimal in (4).

Lemma 5 *In any optimal solution to (6), the constraint binds.*

Proof. Suppose not, and the constraint does not bind at some profile λ which solves the optimization problem. One can alter bidder 1's inverse bid function by making him bid more aggressively (that is, decrease $\lambda_1(b)$) on an interval $[\underline{b}, \bar{b}]$ of bids occurring with positive probability. By continuity, there exists a small perturbation such that the constraint continues to be non-binding. This perturbation increases seller expected revenue, and so the original profile λ was not optimal. ■

Proposition 6 *The maximum of the revenue to the seller among the set of ε -equilibria is attained by a symmetric strategy profile, for ε sufficiently small.*

Proof. Suppose there exists a profile λ that solves the maximization problem (6), but is asymmetric. Without loss of generality, assume that $\lambda_1(b) > \lambda_2(b)$ over an interval of bids $[\underline{b}, \bar{b}]$. Consider a perturbation of λ by two parameters α_1 and α_2 , such that

$$\lambda_1(\alpha_1, \alpha_2) = \lambda_1(\alpha_1) \tag{7}$$

$$\lambda_2(\alpha_1, \alpha_2) = \lambda_2(\alpha_2) \tag{8}$$

$$\lambda_i(\alpha_1, \alpha_2) = \lambda_i, \forall i > 2. \tag{9}$$

That is, λ_1 is a function of b and α_1 only, λ_2 is a function of b and α_2 only, and the remaining bidders' inverse bid functions are held unchanged. Assume that this perturbation is such that

$$\lambda_i(b; 0) = \lambda_i(b) \tag{10}$$

$$\frac{\partial \lambda_1(b; 0)}{\partial \alpha_1} < 0 \tag{11}$$

$$\frac{\partial \lambda_2(b; 0)}{\partial \alpha_2} > 0 \tag{12}$$

$$\frac{\partial \lambda_i(\underline{b}; \alpha)}{\partial \alpha_i} = 0 \tag{13}$$

$$\frac{\partial \lambda_i(\bar{b}; \alpha)}{\partial \alpha_i} = 0 \tag{14}$$

In words, on this interval of bids, move the two bid functions towards each other, making the profile “more symmetric”, in such a way as to hold the set of types being considered unchanged.

The sum on the left hand side of the constraint in problem (6) is made up of three components: The sum of the values of the bidders' best responses, $\sum_{i=1}^n U_i^*(\lambda)$, plus the total surplus extracted by the bidders, $\sum_{i=1}^n S_i(\lambda)$, minus the total amount paid by the bidders to

the seller. To construct the perturbation of λ such that the total payment to the seller is held constant, write $\alpha_2 = a(\alpha_1)$, and note that on the interval $[\underline{b}, \bar{b}]$, the distribution of the maximum bid (and, hence, seller revenue) is kept constant when

$$\frac{\partial \lambda_1(b, \alpha)}{\partial \alpha_1} \lambda_2(b, a(\alpha)) \prod_{i=3}^n \lambda_i(b) + \lambda_1(b, \alpha) \frac{\partial \lambda_2(b, a(\alpha))}{\partial \alpha_2} a'(\alpha) \prod_{i=3}^n \lambda_i(b) = 0 \quad (15)$$

Evaluating at $\alpha = 0$ and rearranging, this yields

$$\frac{\lambda_2(b)}{\lambda_1(b)} = - \frac{\frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} a'(0)}{\frac{\partial \lambda_1(b, 0)}{\partial \alpha_1}} \quad (16)$$

Effectively, this reduces the perturbation of λ to one parameter, which will be written α hereafter.

The surplus extracted by bidder i as a function of α can be written

$$S_i(\alpha) = \int_{\underline{b}}^{\bar{b}} \lambda_i(b, \alpha) \prod_{j \neq i} \lambda_j(b, \alpha) \lambda'_i(b, \alpha) db, \quad (17)$$

where, for brevity, λ'_i refers to the derivative with respect to b . Differentiating with respect to α and setting $\alpha = 0$ yields

$$S'_i(0) = \sum_{j=1}^n \left[\int_{\underline{b}}^{\bar{b}} \frac{\partial \lambda_j(b, 0)}{\partial \alpha_j} \alpha'_j(0) \left(\prod_{k \neq j} \lambda_k \right) \lambda'_i(b) db \right] + \int_{\underline{b}}^{\bar{b}} \frac{\partial^2 \lambda_i(b, 0)}{\partial b \partial \alpha_i} \alpha'_i(0) \left(\prod_{k=1}^n \lambda_k \right) db \quad (18)$$

Rearranging (16) yields

$$\frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} \lambda_2(b) - \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} \alpha'_2(0) \lambda_1(b) = 0 \quad (19)$$

This implies that the summation in (18) equals zero.

For $i > 2$, the second integral in (18) is zero, so

$$\sum_{i=1}^n S'_i(0) = \sum_{i=1}^2 \int_{\underline{b}}^{\bar{b}} \frac{\partial^2 \lambda_i(b, 0)}{\partial b \partial \alpha_i} \alpha'_i(0) \left(\prod_{k=1}^n \lambda_k \right) db \quad (20)$$

Integration by parts yields

$$S'_i(0) = \alpha'_i(0) \lambda_1(b) \lambda_2(b) \frac{\partial \lambda_i(b, 0)}{\partial \alpha_i} \Big|_{b=\underline{b}}^{b=\bar{b}} - \int_{b=\underline{b}}^{b=\bar{b}} \alpha'_i(0) \lambda'_1(b) \lambda_2(b) \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} db \quad (21)$$

$$- \int_{b=\underline{b}}^{b=\bar{b}} \alpha'(0) \lambda_1(b) \lambda'_2(b) \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} db \quad (22)$$

The first term in each equation is zero by assumption, since $\frac{\partial \lambda_i(\bar{b}, \alpha)}{\partial \alpha_i} = \frac{\partial \lambda_i(\underline{b}, \alpha)}{\partial \alpha_i} = 0$. To see that the sum of the second terms is positive, write

$$\int_{b=\underline{b}}^{b=\bar{b}} a'(0) \lambda_1'(b) \lambda_2(b) \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} + \int_{b=\underline{b}}^{b=\bar{b}} \lambda_1'(b) \lambda_2(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} < 0 \quad (23)$$

Since $\lambda_1(b) > \lambda_2(b)$, this is implied by

$$\int_{b=\underline{b}}^{b=\bar{b}} a'(0) \lambda_1'(b) \lambda_1(b) \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} + \int_{b=\underline{b}}^{b=\bar{b}} \lambda_1'(b) \lambda_2(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} \leq 0 \quad (24)$$

Applying the substitution in (16) yields

$$- \int_{b=\underline{b}}^{b=\bar{b}} \frac{\lambda_2(b)}{\lambda_1(b)} \lambda_1'(b) \lambda_1(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} + \int_{b=\underline{b}}^{b=\bar{b}} \lambda_1'(b) \lambda_2(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} \leq 0, \quad (25)$$

which is true.

To see that the sum of the third terms is positive, write

$$\int_{b=\underline{b}}^{b=\bar{b}} a'(0) \lambda_1(b) \lambda_2'(b) \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} + \int_{b=\underline{b}}^{b=\bar{b}} \lambda_1(b) \lambda_2'(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} \leq 0 \quad (26)$$

Applying (16) yields

$$- \int_{b=\underline{b}}^{b=\bar{b}} \frac{\lambda_2(b)}{\lambda_1(b)} \lambda_1(b) \lambda_2'(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} + \int_{b=\underline{b}}^{b=\bar{b}} \lambda_1(b) \lambda_2'(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} \leq 0 \quad (27)$$

which is the same as

$$\int_{b=\underline{b}}^{b=\bar{b}} (\lambda_1(b) - \lambda_2(b)) \lambda_2'(b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} \leq 0 \quad (28)$$

which is true since $\lambda_1(b) > \lambda_2(b)$ and $\frac{\partial \lambda_1}{\partial \alpha_1} < 0$. Thus, the total surplus increases by this perturbation, which has a negative effect on the left hand side of the constraint in (6).

Finally, consider the effect of the perturbation on the values of the best responses. This can be written

$$\sum_{i=1}^n \frac{\partial U_i^*(0)}{\partial \alpha} = \int_{b=\underline{b}}^{b=\bar{b}} (\lambda_1^*(b) - b) \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} a'(0) \lambda_1'(b) db + \int_{b=\underline{b}}^{b=\bar{b}} (\lambda_2^*(b) - b) \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} \lambda_2'(b) db \quad (29)$$

It is enough that this quantity be less negative than the change in efficiency is positive. One condition for this to hold is for λ to be close to the Nash equilibrium. In the neighborhood of the equilibrium,

$$\frac{\partial U_1^*(0)}{\partial \alpha} + \frac{\partial U_2^*(0)}{\partial \alpha} \approx \int_{b=\underline{b}}^{b=\bar{b}} 2b \frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} a'(0) db + \int_{b=\underline{b}}^{b=\bar{b}} 2b \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} db \quad (30)$$

The right hand side is positive if

$$\frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} a'(0) + \frac{\partial \lambda_1(b, 0)}{\partial \alpha_1} > 0 \quad (31)$$

which is true if

$$\frac{\frac{\partial \lambda_2(b, 0)}{\partial \alpha_2} a'(0)}{\frac{\partial \lambda_1(b, 0)}{\partial \alpha_1}} + 1 > 0 \quad (32)$$

which is equivalent to

$$-\frac{\lambda_2(b)}{\lambda_1(b)} + 1 > 0 \quad (33)$$

which is true since $\lambda_1(b) > \lambda_2(b)$. If ε is small, this approximation to the best reply will be valid.

Thus, there exists a perturbation of λ such that the left side of the constraint in (6) decreases strictly, but leaves revenue unchanged. Continuity ensures it is possible to construct a profile that keeps the quantity unchanged and increases seller revenue.

The preceding construction shows that we can take any asymmetric profile feasible in (6) and replace it with another feasible profile which increases seller revenue. Thus, it must be that the maximum seller revenue on the set occurs when the bidders adopt symmetric bid functions. In this case, the magnitudes of the errors made by the bidders are identical, and therefore are all equal to ε . Since the set of ε -equilibria is a subset of the set of profiles feasible in the relaxed problem, these symmetric profiles must also maximize seller revenue on the contained set of ε -equilibria.

The proof for the second-price case is analogous. It is made somewhat easier in that the best reply is always known, and so approximations to the best reply are not necessary. ■

Remark 7 *Since the construction in the proof above uses a set larger than the set of ε -equilibria, it does not rule out the possibility of asymmetric local revenue maximizers within the set of ε -equilibria.*

The intuition behind the proof is that, when asymmetric bid functions are moved “closer together”, the result is to decrease the amount of gains from exchange lost to inefficiency. In some cases, it would be possible that this perturbation of an ε -equilibrium would result in a profile which is not an ε -equilibrium; however, by perturbing in such a way as to hold seller revenue constant, the efficiency gain, essentially, accrues to the bidders, and thus reduces the sum of the optimization errors they make. Therefore, there is a relationship between symmetry and high seller revenues: the seller can be made indirectly better off within the set of ε -equilibria by decreasing the inefficiency which occurs in an asymmetric profile.

Next, I turn to the question of the behavior of the maximum revenue in the region where bidders’ optimization errors are vanishingly small. The following proposition shows that the effect of these tiny errors is substantial.

Proposition 8 *The change in the maximum feasible revenue as a function of ε , $\frac{\partial \bar{R}}{\partial \varepsilon}$, at $\varepsilon = 0$ is infinite.*

Proof. Taking advantage of Proposition 6, write the dual problem to the revenue-maximization problem (4):⁶

$$\begin{aligned} \underline{\varepsilon}(r) = \min_{\lambda \in \Sigma_1} & U(\lambda^*(\lambda), \lambda) - U(\lambda, \lambda) \\ \text{subject to} & R(\lambda, \lambda) \geq r \end{aligned} \quad (34)$$

Write $\lambda(b, \alpha)$ as a perturbation of the equilibrium bid function, where $\lambda(b, 0)$ is the symmetric Nash equilibrium bid function, and $\frac{\partial \lambda(b, 0)}{\partial \alpha} < 0$. For a given perturbation, we can rewrite the problem as

$$\begin{aligned} \underline{\varepsilon}(r) = \min_{\alpha} & U(\alpha^*(\alpha), \alpha) - U(\alpha, \alpha) \\ \text{subject to} & R(\alpha) \geq r \end{aligned} \quad (35)$$

Then, the first-order necessary condition for optimality is given by

$$U^{*'}(\alpha^*(\alpha), \alpha) - \frac{\partial U_1(\alpha)}{\partial \alpha} - \frac{\partial U_2(\alpha)}{\partial \alpha} + \mu \frac{\partial R(\alpha)}{\partial \alpha} = 0, \quad (36)$$

where the subscripts 1 and 2 on U refer to differentiation with respect to the first and second arguments, respectively. At $r = \frac{n-1}{n+1}$, $\alpha = 0$. Recall that, by the envelope theorem, in computing the derivative of the optimal utility with respect to α , the effect of the change in the best reply can be neglected. Thus, $U^{*'}(0, 0) = U_2(0, 0)$. Additionally, in equilibrium, by definition, $U_1(0, 0) = 0$. Since the change in revenue is positive, it must be that $\mu = 0$.

⁶The primal version of the problem fails the constraint qualification at the equilibrium point.

The envelope theorem says that

$$\frac{d\underline{\varepsilon} \left(\frac{n-1}{n+1} \right)}{dr} = \frac{\partial}{\partial r} [U(\lambda^*(\lambda), \lambda) - U(\lambda, \lambda)] - \mu \left(\frac{n-1}{n+1} \right) \frac{\partial}{\partial r} [R(\lambda, \lambda) - r]. \quad (37)$$

The first partial derivative equals zero; the second equals -1, and so

$$\frac{d\underline{\varepsilon} \left(\frac{n-1}{n+1} \right)}{dr} = \mu \left(\frac{n-1}{n+1} \right) = 0, \quad (38)$$

and the claim follows. ■

3.2 Effects on efficiency

The symmetry of the equilibria in these symmetric auctions ensures that the equilibria are efficient; that is, the bidder who values the object most wins it with probability one. Within the set of ε -equilibria, however, there may be substantial asymmetry, depending on the deviations made by the bidders, which will lead to inefficiencies in the allocation of the good.

Efficiency will be measured in terms of the fraction of ex-ante gains from exchange which are realized by an ε -equilibrium.⁷ Formally, the efficiency η is defined as

$$\eta = \frac{\int_{x_1=0}^1 \cdots \int_{x_n=0}^1 x_{\{\arg \max_i b_i(x_i)\}} dx_1 \cdots dx_n}{\int_{x_1=0}^1 \cdots \int_{x_n=0}^1 \max_i x_i dx_1 \cdots dx_n}. \quad (39)$$

Problem 9 (Minimizing percentage of gains from exchange realized)

$$\begin{aligned} \underline{\eta}(\varepsilon) &= \min_{\lambda_1 \in \Sigma_1, \lambda_2 \in \Sigma_2} \eta(\lambda_1, \lambda_2) \\ \text{subject to } & U_i(\lambda_i^*(\lambda_{-i}), \lambda_{-i}) - U_i(\lambda_i, \lambda_{-i}) \leq \varepsilon \quad \forall i \in \{1, \dots, n\} \end{aligned} \quad (40)$$

The minimizing profile for this problem will certainly be asymmetric. As such, characterization of the minimizers is not easily done without resorting to numerics.⁸ Therefore, I will proceed numerically in the next section to determine the amount of inefficiency possible in ε -equilibria in these settings.

⁷An alternate measure of efficiency would be the probability that the bidder who values the object most receives it. Numerical experiments with this measure show that the minimum probability of efficient allocation drops rapidly for increasing ε ; however, as indicated by the calculations reported below, these almost always involve misallocation among bidders with private values close to each other. The gains from exchange measure here also has nicer properties as the number of bidders grows beyond two.

⁸To paraphrase wisdom I once received: “there is only one way to be symmetric, but many ways to be asymmetric”.

4 Numerical results

4.1 Representation of bid functions

To approximate the strategy set of monotonic bid functions, the set of splines of order 4 with $l - 1$ internal knots will be used.⁹ The splines use the uniform grid of “knots” β_k on $[0, \bar{b}_i]$

$$\beta_k = \frac{k}{l} \bar{b}_i, \quad k = 0 \dots l \quad (41)$$

where \bar{b}_i is the maximum bid submitted by bidder i , and assume the inverse bid function is of the form

$$\lambda_i(b) = c_{3,k}^i b^3 + c_{2,k}^i b^2 + c_{1,k}^i b + c_{0,k}^i \text{ on } b \in [\beta_{k-1}, \beta_k]. \quad (42)$$

The coefficients are selected such that these functions are continuous, with continuous first and second derivatives, and $\lambda_i(0) = 0$.¹⁰

Note that the equilibria of both the first- and second-price auctions are exactly representable in this form, and furthermore, remain Nash equilibria in the game where bidding strategies are restricted to the form (42).¹¹ While this form does restrict the possible deviations around the equilibrium point, the set of all deviations can be approximated arbitrarily well by choosing l sufficiently high.

4.2 Algorithm

The constrained optimization problems expressed in equations (4), (5), and (40) need not have a unique local optimum. Additionally, in the case of the efficiency minimization problem, the asymmetries involved ensure that multiple global optima exist, where the profiles are permuted among the bidders. Techniques for solving these problems must take these difficulties into account.

To combat them, I employ a two-phase process to compute the optima. In the first phase, with ε given, a genetic algorithm implements a global search of the set of ε -equilibria. The goal of the genetic algorithm is to organize this search in an intelligent fashion, by taking profiles with relatively large (in the case of minimization, small) values, recombining and

⁹The computations reported used $l = 2$; selected test calculations yielded little difference in the values with l taken as high as 10.

¹⁰See Appendix B for details on why this representation was chosen, and how the splines are parameterized.

¹¹In the second-price auction, since bidding one’s actual type is a dominant strategy, the best reply to any profile is also exactly representable in this family. However, this is not necessarily the case in the first-price auction.

perturbing them. The genetic algorithm approach is relatively immune to getting “stuck” at a local optimum; however, it is also extremely unlikely the algorithm will exactly hit on the point at which the global optimum occurs. Then, given a set of promising starting points as determined by the genetic algorithm, the second phase employs a polytope search algorithm to find the associated local minimum in the region with high precision. This hybrid approach combines the strengths of the two approaches while avoiding their weaknesses.¹²

Note that, inherently, all the estimates for the minima and maxima are really underestimates, in the sense that it is not possible to rule out the existence of a higher value (or, in the case of minimization, lower value) within the set. Ideally, one would like to have some way to construct an overestimate as well, so as to provide a bound on the error of the computation. This is done indirectly by computing the optima for several values of ε , using independently-generated starting populations for each ε . Since the set of ε -equilibria is increasing in ε , the maxima (minima) must be increasing (decreasing); thus, any deviations from this pattern, or unusually sharp changes in the estimates among neighboring ε , would signal that a sufficiently good estimate has not been obtained.

4.3 Effects on seller’s expected revenue

Table 3 in Appendix A gives a list of computed values for $\bar{R}(\varepsilon)$ and $\underline{R}(\varepsilon)$ for selected levels of ε . These are summarized in Figure 1, where the interior pair of lines represents the maximum and minimum of the range of revenues in the first-price auction, and the outer pair of lines plots the range of revenues in the second-price auction. The steepness of the revenue maximum curves near $\varepsilon = 0$ is consistent with Proposition 8.¹³

Numerical Result 10 *For any value of ε in $(0, .001]$, the range of expected revenues to the seller is strictly larger in the second-price auction than in the first price auction.*

Figures 2 and 3 display the profiles which maximize and minimize seller revenue for the first- and second-price auctions, respectively, for the case when $\varepsilon = .001$. The intuition for the different shapes comes from considering the differences in the payment rules. In the first-price auction, as one draws a higher type, one is more likely to be the one setting the price; thus, the bids submitted by these types have a larger effect on seller revenue. In the second-price case, the price-setting bid is more likely to be submitted by a bidder with a

¹²The interested reader can find more specific details on the implementation of the procedure in Appendix C.

¹³A quadratic fit to the computed revenue maximization points yields implied slopes at $\varepsilon = 0$ on the order of 200. Moreover, the R^2 of these fits were both over 0.999, which suggests a conjecture that these curves are in fact quadratic.




Figure 1: Ranges of revenues $[\underline{R}(\varepsilon), \overline{R}(\varepsilon)]$ consistent with ε -equilibrium in first-price (interior range) and second-price (exterior range) auctions, as a function of ε .

type on the low end of the interval of types; thus, more (or less) aggressive bidding in this region drives seller revenue.

As plotted, with the types on the horizontal axes, and bids on the vertical axes, the revenue-maximizing profiles do not appear to be mirror images of the revenue-minimizing ones, when reflected through the Nash profile line. However, looking at the inverse bid functions, by flipping the graphs on their sides, reveals that the vertical distance of the inverse functions from the Nash profile, at each bid, is approximately the same.

Table 1 gives details on the properties of these profiles. To calibrate, recall that in the Nash equilibrium, bidders have an expected payoff of $\frac{1}{6}$, so relative to this value, bidders are making an error of roughly one-half of one percent. However, the range of possible seller revenues extends roughly ten percent on either side of the Nash prediction. Even small optimization errors can lead to large revenue swings.

The result that the range of seller revenues is larger in the second-price auction contrasts with Vickrey's assertion that the second-price auction should be more robust to errors in strategy. This will be discussed in more detail in Section 5.

4.4 Effects on efficiency

The solution of the efficiency minimization problem (40) is undertaken entirely by numeric means. Table 4 in Appendix A lists computed values of $\underline{\eta}(\varepsilon)$ for selected levels of ε . The




Figure 2: $\varepsilon = .001$ -equilibrium profiles of the first-price auction which maximize (top) and minimize (bottom) seller revenue, compared to the Nash profile (center).




Figure 3: $\varepsilon = .001$ -equilibrium profiles of the second-price auction which maximize (top) and minimize (bottom) seller revenue, compared to the Nash profile (center).

Property	First-price max	First-price min	Second-price max	Second-price min
Revenue	0.3784	0.2917	0.3834	0.2873
U_1	0.1440	0.1880	0.1415	0.1897
U_2	0.1444	0.1869	0.1418	0.1897
η	0.99986	0.999993	0.999997	0.999999

Table 1: Economic properties of extreme seller revenue $\varepsilon = .001$ -equilibria for the first- and second-price auctions

results are summarized in Figure 4. From an ex-ante perspective, the deviations of the bidders in these approximate equilibrium have a very small impact on the percentage of surplus lost via inefficiency. However, the first-price auction dominates in terms of avoiding inefficiency in the worst case.




Figure 4: Minimum percentage efficiencies $\underline{\eta}(\varepsilon)$ consistent with ε -equilibrium in first-price (top line) and second-price (bottom line) auctions.

The efficiency-minimizing profile in the second-price auction is particularly interesting. Since high types are likely to win the auction, the utility effect of large changes in bids for these types is very small, a feature which is borne out clearly in the profile. The intertwining of the bid functions for lower types is unusual, and may be an artifact of the functional form being used. Computations with l increased show that the number of crossings appears to be equal to l , leading one to conjecture that the limiting profile is one where bidders bid




Figure 5: $\varepsilon = .001$ -equilibrium profile which minimizes the percentage of efficiency η in the first-price auction.




Figure 6: $\varepsilon = .001$ -equilibrium profile which minimizes the percentage of efficiency η in the second-price auction.

Property	First-price	Second-price
Revenue	0.3294	0.3326
U_1	0.1618	0.1768
U_2	0.1743	0.1528
η	0.9982	0.9930

Table 2: Economic properties of the most inefficient ε -equilibria (in percentage terms) for the first- and second-price auctions

approximately according to their Nash profile for low types, then deviate greatly for high types.

5 Discussion

The conclusions drawn from the computations are in contrast to Vickrey’s assertion that the second-price auction should be more robust to errors in strategy. When bidders are assumed to make small errors in payoff terms, the second-price auction displays both a wider range of revenues to the seller, and a greater possibility of inefficiency via misallocation of the object. As such, they do not support Vickrey’s prediction that the second-price auction would be preferable to all concerned. However, there is experimental evidence that the observed efficiencies in second-price auctions are in fact higher than under the first price rule. Cox et al [4] give data showing the percentage of gains from exchange extracted to be higher in the second-price case.

What drives the implications of ε -equilibrium is the following. The change in expected payoff for a bidder of type x if he bids b in the first-price auction is given by

$$\frac{\partial U}{\partial b} = (x - b) P'(b) - P(b), \quad (43)$$

where $P(b)$ is the probability of winning with bid b . The first term accounts for the change in utility by going from just losing to just winning; the second term accounts for the change in how much is paid in those situations where the bidder would have won anyway.

In the second-price case, the bidder has no effect on price; so, the expression reduces to

$$\frac{\partial U}{\partial b} = (x - b) P'(b). \quad (44)$$

Thus, in the second-price auction, all other things being equal, the derivative is less steep, since the bidder is not disciplined by the effect of his own bid on what he pays. If the

bidders are imagined to be following some iterative method, such as a numerical optimization procedure, to obtain their best reply, they could be fooled by the smaller slope in the second-price auction into stopping the procedure at a bid which is farther from their optimal one.

Herein lies a distinction to be drawn between the concepts underlying dominance, and those underlying an optimization-based method. The dominance criterion asserts that one strategy (in the case of the second-price auction, bidding one's true value) is always no worse than any other. However, it does not address *how much* better it is than the alternatives; thus, dominance may not be robust to perturbations caused by costs imposed by limits to computation.

Merlo and Schotter's distinction between "theorists" and "experimentalists" offers a way of explaining the discrepancy between the predictions of ε -equilibrium and experimental results. In the second-price auction, it is conceivable that a bidder might be able to "chunk" the problem by simplifying the strategy space conceptually into three discrete actions: "bid below my true value", "bid my true value", and "bid above my true value". In doing this, it is easy to see that the first and last options are dominated by the middle option.¹⁴ Thus, this auction environment is one which subjects are easily able to transform into a very simple model, with a solution that is easy to compute globally; in the words of Merlo and Schotter, agents would tend to be "theorists" in this case. The story behind ε -equilibrium assumes that agents are not able to make this type of transformation.

5.1 Critiquing ε -equilibrium

Returning to Radner's discussion of the interpretation of ε , if one were to take literally the idea that the ε in ε -equilibrium is something which agents are actually computing, one is faced with a dilemma. On the one hand, an agent might claim, prior to being presented with a game, that he would be happy to play an ε -equilibrium. But, the only way to know for sure that a given profile is indeed an ε -equilibrium is to compute the best response. If the agent knows his best response, then it is not clear why he would not then choose to play it. If computation is conceptualized as being costly, but switching one's strategy is relatively cheap, then the motives behind an agent's willingness to play an ε -equilibrium become unclear. Thus, ε -equilibrium, taken literally, violates the custom of game theory that a solution concept should make sense both from the modeler's point of view, and the point of view of agents playing the game.

One avenue for rehabilitating ε -equilibrium is not to interpret ε as being a quantity which is computed and known with certainty by the agents. Radner suggested ε could be something

¹⁴In fact, this "chunking" argument is the way in which the dominance argument for this auction was first presented to me.

which is assessed subjectively by the agents in the game. One could imagine agents taking a statistical approach similar to procedures used to estimate the unknown maximum of a distribution.

A more useful interpretation, however, is to think of ε -equilibrium as being a descriptive concept. The definition only speaks to the magnitude of errors in terms of payoff losses, and not to whatever boundedly-rational process is the underlying source of the errors. The observed level of ε appears to depend upon the particulars of the mechanism. Consider, for example, the sealed-bid second-price auction, as described above, versus the ascending-bid implementation of the same auction. It is well-known that these auctions are equivalent; however, experimentally, errors in the ascending-bid version are much smaller than in the sealed-bid version. Thus, the ε in ε -equilibrium is a measure of how much error occurs, and can be used as a guide to designing mechanisms where errors may be smaller.

In this spirit, I now propose two groups of auction settings which are suitable for examining further the set of ε -equilibria, and for experimental determination of the effect of the mechanism on the observed levels of ε .

5.2 Dominance versus Simplicity

The question is, then, it is the dominant-strategy nature of the second-price equilibrium which makes it attractive, or is it the simplicity of that dominant strategy which permit easy solution? To test this, consider an environment in which the second-price auction also has a dominant-strategy equilibrium, but one which is not “simple”, where the dominant nature of the strategy profile is not likely to be obvious. Harstad and Levin [8] present a class of second-price auctions which is dominance-solvable in two steps. The key property of the signal structure in these auctions is that the maximum of the set of signals is a sufficient statistic for the entire collection. An example they give is as follows.

Example 11 *Let S , the common value of the object for sale, be distributed uniformly on $[0, 1]$. Conditional on s , the realized value of S , let each bidder’s private signal X_i be distributed uniformly on $[0, s]$, with the bidders’ signals being independent conditional on s .*

One situation which might fit this model is an auction for oil drilling rights in a tract of land. The common value would be the amount of oil actually present. Each bidder can do a geological test to estimate the amount of oil there; the test is such that, while it may not always show all the oil present, it will never indicate more oil than actually exists.

Since this example fits Harstad and Levin’s class of “maximal attentive” common value auctions, it is dominance solvable. The unique symmetric equilibrium of this second-price

auction, in the two-bidder case, is given by the bid function¹⁵

$$b_i(x_i) = \frac{x_i \ln(x_i)}{x_i - 1}, \quad i = 1, 2. \quad (45)$$

By any reasonable standard, this equilibrium bid function is not “simple”, to the point of containing a transcendental function. Harstad and Levin present maximal attentive auctions as being suitable for use in the laboratory, as they “offer an opportunity to contrast ... the predictive power of Nash equilibrium and dominance criteria”. In this case, they will also allow us to test the usefulness of ε -equilibrium as a solution concept, and thus, in turn, compare the robustness of Nash equilibrium predictions across different auction institutions.

5.3 Third-price auctions

While the maximal attentive auction preserves a flavor of dominance in the second-price case which may permit some disentanglement of simplicity versus dominance, it suffers from the complication of adding common values. The inference problem inherent in bidding in common-values auctions, and the attendant winner’s curse, may well obscure the effects of dominance-solvability.

To this end, an alternate approach is to remain in the independent private values setting, and instead compare the performance of the first- and second-price sealed-bid auctions with that of the third-price auction, as introduced by Kagel and Levin [12].¹⁶ As its name implies, in the third-price auction, the winning bidder pays the amount of the third-highest bid. Kagel and Levin show that, with risk-neutral bidders, the equilibrium bidding strategy is linear in the bidder’s type, and that furthermore, the strategy involves bidding more than the private value.

Intuitively, in the second-price auction, there are relatively easy arguments against bidding either above or below one’s type. In the first-price auction, such an argument exists only against bidding above one’s type; similarly, in the third-price auction, it is easy to argue against bidding less than one’s type. Thus, a comparative experimental design placing the three mechanisms side-by-side allows for pairwise comparison of observed levels of ε , and also permits consideration of the role the nature of the equilibrium itself (in this case, the slope of the bidding function), plays.

The numerical analysis presented above requires knowledge of the bid functions of each bidder. Most experimental auction data are gathered pointwise; that is, each bidder is

¹⁵The unique symmetric equilibrium in the first-price case can be computed numerically from Milgrom and Weber’s characterization.

¹⁶I am especially grateful to Dan Levin for suggesting much of the experimental design I outline here.

given his private value, and is expected to submit a bid given that information. Selten and Buchta [19] have implemented an experiment in which they elicited piecewise-linear bid functions, instead of simply “point bids”, from subjects. Since the equilibria of these three auctions are linear in the private values of the bidders, this procedure can be simplified to eliciting only linear bid functions passing through the origin, that is, essentially, asking only for one number, a slope. In this restricted game, it is easy to compute the best-reply correspondences, and thus measuring ε and studying the learning process in the auction can be simplified greatly.

6 Conclusions

The motivation of this paper is to argue that Nash equilibrium is too sharp a solution concept, and that practical mechanism design should take into account the potential effects of deviations by participants which are small in payoff terms. In the case of independent private values auctions, looking at these deviations yields a surprising prediction favoring the first-price auction over its second-price counterpart; in particular, revenue equivalence and efficiency do not extend to the case of ε -equilibrium.

In designing mechanisms for use in the real world, economists must wear an engineer’s hat. Just as engineers design and test their bridges and buildings to withstand forces well in excess of their intended use, mechanism designers should investigate a given mechanism’s performance under a number of different potential models of agent behavior in an effort to create more robust designs. A potentially interesting line of research is to develop a standard set of such processes, preferably based upon experimental study of human subjects, to be used in a form of stress-testing of mechanisms.

Finally, the computations presented demonstrate the value of numerical work in the analysis of games. The propositions herein were inspired by consideration of the results of the initial numerical calculations, which revealed a connection between efficiency and seller revenue maximization; thus, computation informed the direction of the theoretical work. Numerical analysis also permitted the examination of the asymmetries involved in generating inefficiencies, and in understanding the types of bidding profiles which result in the most extreme deviations from Nash predictions. The complementarities between the two approaches permitted a deeper analysis than would have been possible by either approach alone.

A Computed estimates

ε	First-price	Second-price
.00005	[0.3240, 0.3428]	[0.3227, 0.3442]
.00007	[0.3223, 0.3448]	[0.3211, 0.3460]
.0001	[0.3205, 0.3478]	[0.3185, 0.3486]
.0003	[0.3101, 0.3566]	[0.3082, 0.3600]
.0005	[0.3037, 0.3640]	[0.3006, 0.3679]
.001	[0.2917, 0.3784]	[0.2873, 0.3834]

Table 3: Ranges of expected revenues consistent with ε -equilibria for first- and second-price auctions.

ε	First-price	Second-price
.00005	0.9999	0.9997
.0001	0.99982	0.9994
.0005	0.99911	0.9968
.001	0.9982	0.9930

Table 4: Minimum percentage efficiencies η consistent with ε -equilibria for first- and second-price auctions.

B Computational details

In the computations reported herein, the bid functions were assumed to be strictly monotonic in the bidder's type. Hence, it is possible to write the computations in terms of the inverse bid function, rather than the bid function itself.

To see why this is advantageous, consider bidder 1's expected value integral, which can be written as

$$U_1(b_1) = \int_0^1 (x_1 - b_1(x_1)) \left(\prod_{j=2}^n b_j^{-1}(b_1(x_1)) \right) dx_1, \quad (46)$$

due to the assumed monotonicity of the bid functions.. To evaluate this integral efficiently, it must be possible to invert the bid functions easily. In the case of a bid function approximated by a quadratic (order 3) spline, this could be done in closed form, but at the cost of

involving a square root. The resulting integral is too complicated to evaluate exactly, and as a result relatively expensive numerical methods must be used to obtain accurate values. Given the very small values of ε involved in the calculations, modest errors in the numerical approximation which might be tolerable elsewhere would not be acceptable here.

On the other hand, consider equation (46) rewritten in terms of the inverse bid functions $\lambda_i \equiv b_i^{-1}$:

$$U_1(\lambda_1) = \int_0^{\bar{b}} (\lambda_1(b) - b) \left(\prod_{j=2}^n \lambda_j(b) \right) \lambda_1'(b) db, \quad (47)$$

where \bar{b} is the maximum bid submitted by bidder 1 (that is, $b_1(1)$), and the derivative factor enters due to the change of variables. Now, when approximating the inverse bid functions λ_i by polynomials, the integrands amount to products of polynomials, which can be done in a straightforward and fast manner exactly.¹⁷

In the implementation of the calculations reported, the inverse bid functions were represented by Schumaker's [18] shape-preserving splines. These splines are quadratic (order 3) splines, which are constructed in such a way as to preserve the concavity and convexity properties of the functions being approximated. Preserving concavity and convexity is particularly important in the first-price auction, as the best-reply function depends on the second derivatives of the opponent's bid functions. The inverses were constructed such that $\lambda_i(0) = 0$ and $\lambda_i(\bar{b}_i) = 1$. The parameters of the spline were the $l - 1$ values of the function at the interior knots, plus the value of \bar{b}_i . Within this parameterization, one can represent exactly the inverse bid functions corresponding to the Nash equilibria of both the first- and second-price auctions, and also the best response to any opponent's inverse bid function in the second-price auction.

To derive the best reply in the first-price auction, consider a closed interval of bids $[b_L, b_H]$ such that all the opponents' inverse bid functions are within the same segment of their spline representation; that is the open interval (b_L, b_H) does not contain any knots. On this interval, all opponents' inverse bid functions are continuous, and have continuous first and second derivatives. The first-order condition for the optimal bid of type x_1 of bidder 1 is

$$(x_1 - b) \Lambda(b) - \Lambda'(b) = 0, \quad (48)$$

where $\Lambda(b)$ is the probability of the bid b winning. Taking advantage of the independent and uniform distribution of the opponents' types, this probability is given by the product of

¹⁷The original versions of some of the computations were done using the specification in 46. Switching to the inverse bid function approach resulted in approximately a 500-fold speedup, as 1000 points were being used to accurately estimate the integrals numerically.

the opponents' inverse bid functions at b ,

$$\Lambda(b) = \prod_{j=2}^n \lambda_j(b). \quad (49)$$

The first-order condition can be rewritten entirely in terms of inverse bid functions as

$$(\lambda_1(b) - b) \Lambda(b) - \Lambda'(b) = 0, \quad (50)$$

Since Λ and Λ' are continuous on the interval, differentiating with respect to b and rearranging yields

$$\lambda_1'(b) = 2 - (\lambda_1(b) - b) \frac{\Lambda''(b)}{\Lambda'(b)}. \quad (51)$$

This is a first-order linear differential equation in λ_1 , and since Λ' and Λ'' are continuous on the interval, the equation has a unique solution up to the specification of an initial condition.¹⁸ Since the opponents' inverse bid functions are all continuous, the best reply is continuous. Thus, one proceeds by first solving for the best reply on the first interval (with left endpoint at zero and boundary condition $\lambda_1^*(0) = 0$), and using the value at the right endpoint as the initial condition for the subsequent interval.

If we write the initial condition for an interval as $\lambda_1^*(\beta) = \alpha$, then the best reply on the interval is given by

$$\lambda_1^*(b) = (\alpha - \beta) \frac{\Lambda'(\beta)}{\Lambda'(b)} + \frac{1}{\Lambda'(b)} \{\Lambda(b) - \Lambda(\beta)\} + b \quad (52)$$

Unfortunately, this best reply is generally a rational function. Rational functions can always be integrated exactly in terms of polynomials, rational functions, inverse tangents, and logarithms. Some experimentation with implementing the algorithm to do so by decomposing into partial fractions suggests that a numerical approach to computing the value of the best reply is easiest.

When numerically computing the expected value integral of the best response, it is important to know the maximum bid submitted by bidder 1 in his best response in order to make good use of standard numerical integration procedures, which usually require knowledge of the range of integration. If it were necessary to solve (51) numerically, that information would not be available a priori.

¹⁸Apostol [1], p. 143, Theorem 6.1

C Description of optimization procedures

C.1 Genetic algorithm phase

Begin by randomly generating an initial set $X_1 = \{x^1, \dots, x^m\}$ of points, where each x^i is a direction vector, in the $n(l+2)$ -dimensional space of coefficients, relative to the equilibrium point \mathbf{a} . Follow this ray from the equilibrium point out to the nearest point such that the ε associated with the profile $\mathbf{a} + \alpha x^i$ is equal to the given ε , for some α . Along this line, compute the maximum value of the seller's expected revenue¹⁹, which will be denoted by ρ^i .

Each direction x^i is assigned a fitness score equal to the amount by which the revenue exceed the equilibrium value,²⁰

$$\phi(x^i) = \max\left\{\rho^i - \frac{1}{3}, 0\right\}. \quad (54)$$

For each point x^i , compute the probabilities

$$p_i = \frac{\phi(x^i)}{\sum_{x \in X_k} \phi(x)}. \quad (55)$$

These can be interpreted as the probability that point x^i reproduces in the next generation. Next, randomly select new populations Y_k of directions out of X_k by sampling m points from X_k , with replacement. From the set Y_k , randomly make $\frac{m}{2}$ pairs of points; denote a typical pair by $\{y^1, y^2\}$. For each coordinate $j \in \{1, 2, \dots, n(l+2)\}$, with probability π , swap the values of y_j^1 and y_j^2 . Define the set Z_k to be the points resulting from this “mating” process.

Finally, with probability q , independently for each point z^i and each dimension, “mutate” the points by adding a random value with mean zero to that dimension of that point. Call the resulting set of directions X_{k+1} , and repeat.

C.2 Polytope search phase

Let $f(x)$ be the quantity being minimized, and start with a simplex in $R^{n(l+2)}$, with vertices $\{x^1, \dots, x^{n(l+2)+1}\}$. Order the vertices such that

$$f(x^i) \geq f(x^{i+1}) \quad \forall i \in \{1, \dots, n(l+2)\}. \quad (56)$$

¹⁹The procedure is analogous for the other quantities of interest.

²⁰In the case of minimization, the fitness score is

$$\phi(x^i) = \max\left\{-\rho^i + \frac{1}{3}, 0\right\}. \quad (53)$$

For other quantities, the definition is analogous.

Begin with x^1 , and reflect it through the opposing face to y^1 . If $f(y^1) < f(x^1)$, replace x^1 with y^1 , reorder the points such that (56) holds, and repeat. Otherwise, proceed to do the same with points $i = 2, \dots, n(l+2)+1$. If all these reflections fail to yield a smaller value, shrink the simplex toward the point $x^{n(l+2)+1}$, and repeat. Once the simplex shrinks to a predetermined size, stop.

References

- [1] Apostol, T. M. *Calculus: Volume II*. New York: John Wiley & Sons, 1969.
- [2] Bajari, P. Auction models are not robust when bidders make small mistakes. Mimeo, 1999.
- [3] de Boor, C. *A Practical Guide to Splines*. New York: Springer, 1978.
- [4] Cox, J. C., B. Roberson and J. M. Walker. Theory and behavior of single object auctions. In *Research in Experimental Economics*, volume 2, V. L. Smith, ed. Greenwich CT: JAI Press, 1-43, 1982.
- [5] Cox, J. C., V. L. Smith and J. M. Walker. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty*, **1**, 61-99, 1988.
- [6] Cox, J. C., V. L. Smith and J. M. Walker. Theory and misbehavior of first-price auctions: Comment. *American Economic Review*, **82**, 1392-1412, 1992.
- [7] Davis, D. D. and C. A. Holt. *Experimental Economics*. Princeton: Princeton University Press, 1993.
- [8] Harstad, R. M. and Dan Levin. A class of dominance solvable common-value auctions. *Review of Economic Studies*, **52**: 525-528, 1985.
- [9] Harrison, G. W. Theory and misbehavior of first-price auctions. *American Economic Review*, **79**, 749-62, 1989.
- [10] Harrison, G. W. Theory and misbehavior of first-price auctions: Reply. *American Economic Review*, **82**, 1426-43, 1992.
- [11] Judd, K. L. *Numerical Methods in Economics*. Cambridge, MA: The MIT Press, 1998.

- [12] Kagel, J. H. and D. Levin. Independent private value auctions: bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. *Economic Journal*, **103**(419), 868-78, 1993.
- [13] Kagel, J. H. and A. E. Roth. Theory and misbehavior of first-price auctions: Comment. *American Economic Review*, **82**: 1379-91, 1992.
- [14] McKelvey, R. D. and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, **10**, 6-38, 1995.
- [15] Merlo, A. and A. Schotter. Theory and misbehavior of first-price auctions: Comment. *American Economic Review*, **82**: 1413-25, 1992.
- [16] Milgrom, P. and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica*, **50**: 1089-1122, 1982.
- [17] Radner, R. Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, **22**: 136-156, 1980.
- [18] Schumaker, L. L. On shape-preserving quadratic spline interpolation. *SIAM Journal of Numerical Analysis*, **20**: 854-64, 1983.
- [19] Selten, R. and J. Buchta. Experimental sealed bid first price auctions with directly observed bid functions. Universitat Bonn Sonderforschungsbereich 303, Discussion Paper B-270, February 1994.
- [20] Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, **16**: 8-37, 1961.