

BEHAVIORALLY OPTIMAL AUCTION DESIGN: EXAMPLES AND OBSERVATIONS

Vincent P. Crawford
University of California, San Diego

Zvika Neeman
Boston University

Tamar Kugler
University of Arizona

Ady Pauzner
Tel Aviv University

Abstract

This paper begins to explore behavioral mechanism design, replacing equilibrium by a model based on “level- k ” thinking, which has strong support in experiments. In representative examples, we consider optimal sealed-bid auctions with two symmetric bidders who have independent private values, assuming that the designer knows the distribution of level- k bidders. We show that in a first-price auction, level- k bidding changes the optimal reserve price and often yields expected revenue that exceeds Myerson’s (1981) bound; and that an exotic auction that exploits bidders’ non-equilibrium beliefs can far exceed the revenue bound. We close with some general observations about level- k auction design. (JEL: C72, C92)

1. Introduction

A number of recent papers reconsider core microeconomic questions taking a “behavioral” view of individual decisions or probabilistic judgment (e.g., Camerer, Loewenstein, and Rabin 2003). Most such papers focus on consumer behavior, but a few analyze questions in mechanism design (Glazer and Rubinstein 1998; Neeman 2003; Eliaz and Spiegel 2006, 2007). In those papers, however, the behavioral aspect is limited to decisions or judgment rather than beliefs: Despite the central role of equilibrium assumptions in the theory of mechanism design, there are very few analyses of design outside the equilibrium paradigm.¹

Taking a broader view of strategic behavior should increase the practical usefulness of mechanism design theory. Design inherently involves the creation of

Acknowledgments: Thanks to Eddie Dekel, Yehoshua Hoffer, Nagore Iriberry, Navin Kartik, Juanjuan Meng, Roger Myerson, Robert Östling, Adam Sanjurjo, and Joel Sobel for helpful discussions.

E-mail addresses: Crawford: vcrawfor@dss.ucsd.edu; Kugler: tkugler@eller.arizona.edu; Neeman: zvika@bu.edu; Pauzner: pauzner@post.tau.ac.il

1. Matsushima (2007, 2008) takes a finitely iterated dominance approach similar in spirit to but quite different in substance from the one we take here.

new games, for which the learning justification for equilibrium may be weak or nonexistent; yet it is often important for an application to work the first time.² Further, assuming equilibrium may yield theoretically optimal designs that are too complex for confidence in equilibrium behavior, even if learning is possible. Replacing equilibrium with a model that better describes people's responses to new and/or complex games—both equilibrium responses in games that tend to elicit them and systematic deviations in games that don't—should allow us to design more effective mechanisms. It also suggests a concrete, evidence-based way to assess robustness, something previously left to intuition.

This paper begins to explore relaxing the equilibrium assumption in mechanism design. We consider the leading case of an optimal (expected-revenue maximizing) single-object sealed-bid auction with two symmetric bidders who have independent private values, for which Myerson (1981) provides a complete equilibrium-based analysis.

To focus sharply on strategic behavior, we maintain the standard rationality assumptions regarding decisions and judgment. We model strategic behavior via a class of structural non-equilibrium models called level- k models, which have strong experimental support from experiments that elicit initial responses to a variety of games.³ In a level- k model players' decision rules or "types" are allowed to be heterogeneous, but each player's type is drawn from a common distribution. Type Lk anchors its beliefs in a nonstrategic $L0$ type, which represents its model of others' instinctive reactions to the game as discussed subsequently, and adjusts them via thought-experiments with iterated best responses: $L1$ best responds to $L0$, $L2$ to $L1$, and so on.⁴ $L1$ and higher types are rational, with perfect models of the game. Their only departure from equilibrium is replacing its perfect model of others with a simplified model. In applications the type frequencies are estimated or calibrated from previous analyses. The estimated $L0$ frequency is usually small; and most weight is on $L1$ and $L2$.

Like equilibrium, level- k models are general, tractable, and make specific predictions. Unlike equilibrium, their predictions are stochastic and depend on a small number of behavioral parameters—the type frequencies, and precisions

2. In the best-known such example, the U.S. FCC spectrum auction (McMillan 1994), billions of dollars were at stake. Partly because they believed sealed-bid designs would not yield outcomes close enough to equilibrium to ensure good results, the designers adopted a progressive, partly "open" design for which the theory was weaker but experimental results were more promising.

3. Costa-Gomes and Crawford (2006) and Costa-Gomes, Crawford, and Iriberri (2009) review the literature on level- k models. Our level- k model most closely follows that of Crawford and Iriberri (2007), who showed how to adapt level- k models to auctions with incomplete information.

4. Lk respects k rounds of iterated dominance, so a distribution of low- k types mimics equilibrium in games that are dominance-solvable in few rounds, but may deviate in more complex games, in predictable ways. Thus a "robust" mechanism that implements good outcomes and is dominance-solvable in one or two rounds (compare Matsushima 2007, 2008) may have an advantage over a complex mechanism that theoretically implements better outcomes but only in equilibrium.

that describe the noisiness of types' decisions—but they have enough structure to allow an informative analysis, contingent on realistic parameter values.⁵

We assume that bidders are drawn from a given population of level- k types, whose frequencies are known to the designer. Because the question of optimal auctions with level- k bidders is a difficult one, most of our analysis considers representative examples and the narrower questions of what reserve prices are optimal and how much revenue they yield in first-price auctions, the optimality of auction forms, and the use of exotic auctions that exploit bidders' non-equilibrium beliefs to exceed Myerson's (1981) revenue bound.

Section 2 introduces the main examples we study and reviews Myerson's equilibrium analysis of optimal auctions, following Krishna (2002). Section 3 defines our level- k model and develops its implications for revenue-maximizing reserve prices in first-price auctions in the main examples, showing that the designer can often adjust the reserve to do better than the maximal revenue with equilibrium bidders. Section 4 gives a different example, in which with equilibrium bidders the optimal auction is a second-price auction with reserve price, but in which an exotic auction can theoretically (if bidders' level- k behavioral rules are assumed invariant to the design) exploit bidders' non-equilibrium beliefs to obtain expected revenues that far exceed the maximal revenue assuming equilibrium. Section 5 concludes with general observations on level- k auction design, noting problems that arise if the design affects the form of bidders' strategic decision rules and considering difficulties in adapting the revelation principle to characterize optimal auctions with level- k bidders.

2. Equilibrium Analysis of Optimal Auctions

We consider single-object auctions with two risk-neutral bidders whose values are independently and identically distributed (i.i.d.). In this section and the next we consider two examples, one with increasing and one with decreasing value density, which lead to different patterns of level- k deviations from equilibrium that, taken together, are representative of the possibilities.⁶

In the increasing (I) example values have the distribution function $F_{I\gamma}(v) = v^\gamma$ on $[0, 1]$, with the density $f_{I\gamma}(v) = \gamma v^{\gamma-1}$. $\gamma > 0$ is required for $F_{I\gamma}(v)$

5. Costa-Gomes and Crawford (2006) and Costa-Gomes et al. (2009) summarize the evidence in support of level- k models and present the case for using them to model initial responses to games. Until recently the choices for modeling nonequilibrium behavior were limited to rationalizability; adding noise to equilibrium predictions as in Palfrey (2000); or using McKelvey and Palfrey's (1995) notion of quantal response equilibrium (QRE). To our knowledge these models have not been applied to design, perhaps because rationalizability is imprecise (but see Battigalli and Siniscalchi 2003); QRE must be solved numerically; and in noisy equilibrium or QRE the error specification is crucial but has little to guide it (Costa-Gomes, Crawford, and Iriberry 2009).

6. These examples are the basis of the design for experiments we are conducting to guide our specification and test the ideas presented here. The results will be reported in a companion paper.

to be a valid distribution function, and we strengthen this to $\gamma > 1$ to make the density increasing. We suppress γ , writing F_I and f_I instead of $F_{I\gamma}$ and $f_{I\gamma}$. Because

$$v - \frac{1 - F_I(v)}{f_I(v)} = v - \frac{1 - v^\gamma}{\gamma v^{\gamma-1}} = \frac{(\gamma + 1)v}{\gamma} - \frac{1}{\gamma v^{\gamma-1}}$$

is increasing in v when $\gamma > 1$, F_I is “regular” in Myerson’s (1981) sense. Thus, Myerson’s famous result establishes that among the optimal mechanisms in this environment is a first-price auction with suitably chosen reserve price.

In a first-price auction with reserve price $r \in [0, 1]$, the equilibrium bid for value $v \geq r$ can be shown to be (omitted calculations here and below are in an appendix at <http://dss.ucsd.edu/~vcrawfor/#BOAD>)

$$b_{I|r}^E(v) = \frac{r^{\gamma+1}}{(\gamma + 1)v^\gamma} + \frac{\gamma}{\gamma + 1}v,$$

which increases from r at $v = r$ to

$$\frac{r^{\gamma+1}}{\gamma + 1} + \frac{\gamma}{\gamma + 1}$$

at $v = 1$. The optimal reserve price $r_I^{*E}(\gamma)$ can be shown to be

$$r_I^{*E}(\gamma) = (\gamma + 1)^{-\frac{1}{\gamma}}.$$

In the decreasing (D) example values have the distribution function

$$F_{D\alpha\beta}(v) = \frac{\beta}{\beta - \alpha} \left(1 - \frac{\alpha}{v}\right) \text{ on } [\alpha, \beta], \text{ with } \alpha > 0,$$

with well-defined density $\alpha\beta/[(\beta - \alpha)v^2]$ that is positive, decreasing, and continuous on $[\alpha, \beta]$. We suppress α and β , writing F_D and f_D instead of $F_{D\alpha\beta}$ and $f_{D\alpha\beta}$. Because the virtual valuation

$$v - \frac{1 - F_D(r)}{f_D(v)} = \frac{v^2}{\beta}$$

is increasing in v on $[\alpha, \beta]$, F_D is regular. Thus, Myerson’s (1981) result again establishes that an optimal mechanism in this environment is a first-price auction with reserve price. Moreover, because the virtual valuation is always positive, the optimal reserve price must be set to the minimum value, that is,

$$r_D^{*E}(\alpha, \beta) = \alpha.$$

In a first-price auction with reserve price $r \in [\alpha, \beta]$, the equilibrium bid for value $v \geq \max\{\alpha, r\}$ is (see, e.g., Krishna 2002, p. 24)

$$b_{D|r}^E(v) = r \frac{F_D(r)}{F_D(v)} + \frac{1}{F_D(v)} \int_r^v x f_D(x) dx = \frac{v}{v - \alpha} [r - \alpha + \alpha(\ln v - \ln r)],$$

with obvious simplifications when r takes its optimal value of α .

Note that for this distribution, $b_{D|r}^E(v)$ is independent of β . Further, although $b_{D|r}^E(v) \rightarrow \infty$ as $v \rightarrow \infty$, it increases extremely slowly, like $\ln v$. As a result, if $r = \alpha = 1$, for example, $b_{D|r}^E(1,000,000) = 13.82$, so that in equilibrium, a bidder who values the object at \$1,000,000 bids only \$13.82.

It follows that the expected revenue to the seller under the optimal auction is

$$\Pi_{D\alpha\beta|r=\alpha} = \alpha + \frac{\alpha^2}{(\beta - \alpha)^2} \int_{\alpha}^{\beta} \frac{(\beta - v)^2}{v^2} dv.$$

$\Pi_{D\alpha\beta|r=\alpha}$ is bounded above by $2\alpha\beta/(\beta - \alpha) \approx 2\alpha$ for $\beta \gg \alpha$. We use this fact subsequently when we compare the seller's equilibrium and level- k expected revenues.

3. Level- k Analysis of Optimal Reserve Prices in First-Price Auctions

Recall that in a level- k model, type Lk anchors its beliefs in an $L0$ type and adjusts them via iterated best responses: $L1$ best responds to $L0$, $L2$ to $L1$, and so on. To complete the specification, we define the $L0$ type following Crawford and Iriberri (2007), who found that in a variety of auction experiments subjects' initial bids could be described as a mixture of $L1$ or $L2$ responses to either a "random" $L0$ that bids uniformly over the natural range of bids (as in most previous level- k analyses) or a "truthful" $L0$ that bids its private value. Like Crawford and Iriberri, we call the associated $L1$ s or $L2$ s "random" or "truthful" $L1$ s or $L2$ s even though they are not themselves random or truthful.⁷

In our analysis (unlike Crawford and Iriberri's), reserve prices r above the lowest possible value are potentially important, and this creates an ambiguity regarding the "natural range of bids," which for the I example could be either $[0, 1]$ or (truncating the value distribution) $[r, 1]$; and for the D example could be either $[\alpha, \beta]$ or $[r, \beta]$. We focus here on the latter specifications, which our pilot experiments suggest are more descriptive of most subjects' bidding behavior.

7. Using initial-response data from common- or independent-private-value auction experiments, Crawford and Iriberri estimated large frequencies (59–65%) of random $L1$ bidders (in auctions without reserve prices, hence not distinguishing the alternatives below) and much smaller but significant frequencies of random $L2$ (4–9%), truthful $L1$ (9–18%), and truthful $L2$ (1–16%).

Given this specification, for the I example with $r \in [0, 1]$, random or truthful $L1$ s and $L2$ s either decline to bid or, equivalently, bid less than r when $v < r$. When $v \geq r$, level- k types bid as follows:

Random $L1$'s bid is given by

$$b_{I|r}^{1R}(v) \in \arg \max_{b \geq r} (v - b) \frac{b - r}{1 - r} = \frac{r + v}{2},$$

which can be either greater or less than $b_{I|r}^E(v)$.

Truthful $L1$'s bid is given by

$$b_{I|r}^{1T}(v) \in \arg \max_{b \geq r} (v - b)b^\gamma = \max \left\{ r, \frac{\gamma}{\gamma + 1}v \right\} \leq b_{I|r}^E(v).$$

(Here and subsequently, the inequality becomes strict when $v > r$.)

Random $L2$'s bid is given by

$$b_{I|r}^{2R}(v) \in \arg \max_{b \geq r} (v - b)(2b - r)^\gamma = \max \left\{ r, \frac{r}{\gamma + 1} + \frac{\gamma - 1}{\gamma + 1}v \right\} \leq b_{I|r}^E(v).$$

Truthful $L2$'s bid is the same as truthful $L1$'s

$$b_{I|r}^{2T}(v) = \arg \max_{b \geq r} (v - b) \left(\frac{\gamma + 1}{\gamma} \right)^\gamma b^\gamma = \max \left\{ r, \frac{\gamma}{\gamma + 1}v \right\} \leq b_{I|r}^E(v).$$

Thus in the I example, all types except possibly random $L1$ underbid, relative to equilibrium. Examples suggest that underbidding predominates even for random $L1$. Intuitively, random types tend to underbid here because they anchor their beliefs in a uniform random $L0$, which makes them respond to the increasing density less than equilibrium bidders do.

In characterizing the optimal reserve prices for level- k bidders, we focus on the cases in which both bidders are the same type. In the I example, the optimal reserve price facing two random $L1$ bidders, $r_I^{*1R}(\gamma)$, can be shown to be

$$r_I^{*1R}(\gamma) = (4\gamma + 1)^{-\frac{1}{2\gamma}},$$

which is greater (less) than the optimal equilibrium reserve price $r_I^{*E}(\gamma)$ if $\gamma < (>)2$. We have been unable to solve for the optimal reserve prices facing two random $L2$ bidders, two truthful $L1$ bidders, or two truthful $L2$ bidders; but it can be shown numerically that that they can be either higher or lower than the optimal equilibrium reserve price. For low γ and sufficiently high r , random $L2$, truthful $L1$, and truthful $L2$ all bid r for all values v . As a result, the optimal reserve price for low γ for these three types are identical.

In the D example with $r \in [\alpha, \beta]$, random or truthful $L1$ s and $L2$ s either decline to bid or bid less than r when $v < r$. When $v \geq r$, they bid as follows. Random $L1$'s bid is given by

$$b_{D|r}^{1R}(v) \in \arg \max_{r \leq b \leq \beta} (v - b) \frac{b - r}{\beta - r} = \frac{v + r}{2} \geq b_{D|r}^E(v).$$

(Here and in the following, the inequality becomes strict when $v > r$.)

Truthful $L1$'s bid is given by

$$b_{D|r}^{1T}(v) \in \arg \max_{r \leq b \leq \beta} (v - b)F(b) = \max \left\{ r, (\alpha v)^{1/2} \right\} \geq b_{D|r}^E(v).$$

Random $L2$'s bid is given by

$$\begin{aligned} b_{D|r}^{2R}(v) &\in \arg \max_{r \leq b \leq \beta} (v - b)F((b_{D|r}^{1R})^{-1}(b)) \\ &= \arg \max (v - b)F(\max\{r, 2b - r\}) = \frac{r}{2} + \frac{\alpha}{2} \sqrt{\frac{2v - r}{\alpha}} \geq b_{D|r}^E(v). \end{aligned}$$

Finally, truthful $L2$'s bid is given by

$$\begin{aligned} b_{D|r}^{2T}(v) &= \arg \max_{r \leq b \leq \beta} (v - b)F((b_{D|r}^{1T})^{-1}(b)) \\ &= \arg \max_{r \leq b \leq \beta} (v - b)[\beta/(\beta - \alpha)][1 - \alpha^2/b^2] \\ &= \max\{r, z - \alpha^2/3z\}, \end{aligned}$$

where $z = [(\alpha^4 v^2 + \alpha^6/27)^{1/2} + \alpha^2 v]^{1/3}$, so that $b_{D|r}^{2T}(v) \geq b_{D|r}^E(v)$.

Thus in the D example, all types overbid relative to equilibrium. Intuitively, random types tend to overbid with declining densities because they anchor their beliefs in a uniform random $L0$, whereas equilibrium bidders take into account the fact that higher values in the support are less likely than lower values.

Note that for large values of v , $b_{D|r}^{1R}(v)$ is approximately linear in v , $b_{D|r}^{1T}(v)$ and $b_{D|r}^{2R}(v)$ are proportional to $v^{1/2}$, and $b_{D|r}^{2T}(v)$ is proportional to $v^{1/3}$. By contrast, $b_{D|r}^E(v)$ is proportional to $\ln v$, which is much smaller for large values of v . For example, if $r = \alpha = 1$, then $b_{D|r}^{1R}(1,000,000) = 500,000.5$ and $b_{D|r}^{1T}(1,000,000) = 1,000$, much larger than $b_{D|r}^E(1,000,000) = 13.82$.

In characterizing optimal level- k reserve prices in the D example, we again take both bidders to be of the same type. It can then be shown numerically that the optimal reserve prices facing two random $L1$ or two random $L2$ bidders are both higher than the optimal equilibrium reserve price, $r_D^{*E}(\alpha, \beta) = \alpha$. The optimal

reserve prices facing two truthful $L1$ or two truthful $L2$ bidders are both equal to the optimal equilibrium reserve price.⁸

The optimal reserve prices for random $L1$ or $L2$ bidders in the D example illustrate an interesting and potentially important point. It is a commonplace that price lists in other settings create “menu effects,” whereby the presence of high-priced items increases buyers’ demands for lower-priced items (Kantor 2006). The reserve price can evoke a similar effect in an auction, but in the D example with equilibrium bidders this effect is not strong enough to make a reserve price a useful instrument. But in the D example with random Lk bidders the reserve price has a stronger effect, which arises because Lk bidders anchor on an $L0$ that is uniform random on $[r, \beta]$. As a result, a reserve price becomes a useful instrument, although the optimal reserve price remains small.⁹

Given that in the D example, all level- k types bid more aggressively than equilibrium bidders, a designer facing a known distribution of level- k bidders should be able to realize more expected revenue than is possible with equilibrium bidders. We can derive a lower bound on this revenue by calculating it for two random $L1$ bidders and reserve price $r = 0$ and then multiplying by the probability that two such bidders are drawn. It is feasible for the designer to design optimally for this contingency, ignoring all others, and the optimal design can do no worse. Setting $G(v) = F_D(v)^2$, the expected revenue from two random $L1$ bidders and reserve price $r = 0$ is

$$\int_{\alpha}^{\beta} \frac{v}{2} dG(v) = \frac{\alpha\beta^2}{(\beta - \alpha)^2} \left(\ln \beta - \ln \alpha + \frac{\alpha}{\beta} - 1 \right).$$

Notably, this expression $\rightarrow \infty$ (albeit slowly, like $\ln \beta$) as α is held fixed and $\beta \rightarrow \infty$. If $\alpha = 1$ and $\beta = 5,000$, the value of this expression is approximately 7.5, almost four times as large as the 2α that approximates the seller’s expected revenue with equilibrium bidders. Because the probability of drawing two random $L1$ s is about $1/3$, it is clear that in at least some cases, the seller can realize more expected revenue than with equilibrium bidders.

As noted in Section 2, for the D example, a second-price auction with reserve price $r \leq \alpha$ is optimal with equilibrium bidders. But because (with independent private values) a second-price auction makes the equilibrium bid a dominant strategy, level- k bids coincide with equilibrium bids. Thus, with level- k bidders a second-price auction yields only the equilibrium expected revenue. Our analysis therefore shows that in the D example with level- k bidders, a first-price auction

8. For random $L1$ or random $L2$ bidders, the revenue function is single-peaked at some $r > \alpha$. For truthful $L1$ or truthful $L2$ bidders, it declines steadily over the entire range of values $[\alpha, \beta]$.

9. For example, in the D example with $\alpha = 1$ and $\beta = 10$, the optimal reserve price is ≈ 1.33 for random $L1$ bidders and ≈ 1.25 for random $L2$ s, slightly above the lowest possible value.

with suitable reserve price yields higher expected revenue than the best second-price auction. Unsurprisingly, this also shows that revenue-equivalence breaks down with level- k bidders. Further, the advantage of using a first-price auction relative to the simpler and perhaps cheaper method of selling at a posted price is larger with level- k than with equilibrium bidders.

As we have seen, the optimal reserve price can be large with equilibrium bidders but small with level- k bidders, or vice versa. It would be interesting to find general conditions that determine when a reserve price induces more aggressive bidding for equilibrium than level- k bidders, and the extent to which this makes optimal level- k higher than optimal equilibrium reserve prices.

4. Exotic Auctions that Exploit Level- k Bidders' Non-Equilibrium Beliefs

We now give an example in a slightly different environment that illustrates the fact that a designer can exploit level- k bidders' non-equilibrium beliefs to obtain very large expected revenues. As in Sections 2 and 3, we consider a single-object auction with two risk-neutral bidders whose values are independently and identically distributed, but we now suppose that the values are uniformly distributed on the unit interval. The maximum expected surplus (ignoring incentive constraints) for this environment is $E[\max\{v_1, v_2\}] = 2/3$. Myerson (1981) showed that a second-price auction with reserve price 0.5 is optimal here, and it can be shown to yield expected revenue ≈ 0.417 .

Consider the following exotic auction. Bidders submit simultaneous sealed bids $b_1, b_2 \in [0, 1]$. A bidder who bids 1 wins the object if the other bids less than 1. If both bid 1, the winning bidder is chosen randomly. A bidder who bids 1 pays 0 if the other bidder bids less than 1 and pays $M > 1$ if the other bidder bids 1. A bidder who bids less than 1 pays nothing, but cannot win the object.

For this auction, with the uniform value distribution, both truthful and random $L0$ bidders imply uniform random bid distributions on the unit interval. As a result, truthful and random $L1$, defined as in Section 3 for an auction with no reserve, both believe that if they bid 1 they will win the object and pay 0. Given truthful and random $L0$ s' bid distributions, truthful and random $L1$ assign zero prior probability to the possibility that they will have to pay M if they bid 1. Consequently, truthful and random $L1$ are willing to participate, and both bid 1, whatever their values. Given this behavior, both truthful and random $L2$ believe that the other bidder will bid 1 for certain, and therefore expect that bidding 1 will result in their winning the object for the price of $M > 1$ with probability 0.5. Hence they decline to bid or, equivalently, bid less than 1.

Truthful and random $L3, L5, \dots$, behave the same as truthful and random $L1$; and truthful and random $L4, L6, \dots$, behave the same as truthful and random $L2$. Thus when both bidders have k odd the seller's expected revenue is M , and

when at least one of the bidders has k even it is 0. Even a designer who does not know the type distribution can obtain very large expected revenue by setting M very large to exploit level- k bidders' non-equilibrium beliefs.

5. General Observations on Level- k Auction Design

This section concludes with some general observations on level- k auction design. First, our formulation of the design problem takes the level- k model's specification as given, independent of the design, just as the standard analysis assumes that bidders will play an equilibrium for any design. Although our specification is based on substantial experimental evidence and is general enough to apply to any game, there is reason to doubt this assumption for exotic auctions like Section 4's, which go beyond the evidence on which our level- k specification is based (Crawford and Iriberry 2007). For example, bidders might view Section 4's auction as having a reserve price of 1, in which case the most natural $L0$ specification (random or truthful) implies a spike of bids at 1, with the result that a value of M high enough to be profitable would make even truthful and random $L1$ bidders decline to bid. More generally, an auction design that is optimal for a level- k specification when it is assumed independent of the design might not be an auction for which the level- k model describes behavior well.¹⁰ A general formulation of the design problem must take a position on how the design influences the rules that describe bidders' behavior and develop new methods to deal mathematically with that influence.

Even without influences like those just discussed, the heterogeneity of level- k strategic decision rules and beliefs complicates the characterization of optimal auctions.¹¹ In the standard analysis there is no loss of generality in using the revelation principle to restrict attention to incentive-compatible direct mechanisms because if equilibrium is assumed, each bidder can check that all bidders have an incentive to reveal their values truthfully. Given the restriction to direct mechanisms, the design problem is guaranteed to have a solution.

By contrast, Section 4's example shows that even with a known, homogeneous population of level- k bidders, the differences between their beliefs about their own and others' behavior may be exploitable in a way that makes the design problem as usually formulated fail to have a solution. There are also problems with extending the revelation principle to this setting, which we take to mean finding a direct mechanism that creates incentives for all bidders to reveal their private values, while respecting (in the associated indirect mechanism) their wishes to follow

10. A similar criticism applies to equilibrium auction design. For example, a first-price auction with no reserve price is theoretically optimal in our D example with equilibrium bidders, but we believe very few actual bidders with values of \$1,000,000 would bid as low as \$13.82.

11. This is somewhat surprising because Crawford and Iriberry (2007) show that many other results from equilibrium auction theory generalize straightforwardly to level- k auction theory.

their non-equilibrium decision rules. For example, a direct mechanism that would give a bidder an incentive to reveal his value truthfully, assuming that other bidders reveal truthfully, might not convince the bidder that others will reveal truthfully, because his model of others' decision rules is different from his own rule. Our preliminary analyses suggest that there are important cases where progress can nonetheless be made: If, for example, all bidders are known to be truthful L_k for some k , the incentive constraint for L_1 s also ensures truthful revelation for higher L_k s, and L_0 is automatically truthful. In this case, and possibly some others, it may be possible to adapt Myerson's (1981) equilibrium methods to characterize optimal auctions for level- k bidders. Such analyses pose interesting challenges, which should increase the practical usefulness of design theory.

References

- Battigalli, Pierpaolo, and Marciano Siniscalchi (2003). "Rationalizable Bidding in First-Price Auctions." *Games and Economic Behavior*, 45, 38–72.
- Camerer, Colin, George Loewenstein, and Matthew Rabin (2003). *Advances in Behavioral Economics*. Princeton University Press.
- Costa-Gomes, Miguel, and Vincent Crawford (2006). "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review*, 96, 1737–1768.
- Costa-Gomes, Miguel, Vincent Crawford, and Nagore Iriberry (2009). "Comparing Models of Strategic Thinking in Van Huyck, Battalio, and Beil's Coordination Games." *Journal of the European Economic Association Papers and Proceedings*, 7, 365–376.
- Crawford, Vincent, and Nagore Iriberry (2007). "Level- k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 75, 1721–1770.
- Eliasz, Kfir, and Ran Spiegler (2006). "Contracting with Diversely Naïve Agents." *Review of Economic Studies*, 73, 689–714.
- Eliasz, Kfir, and Ran Spiegler (2007). "A Mechanism-Design Approach to Speculative Trade." *Econometrica*, 75, 875–884.
- Glazer, Jacob, and Ariel Rubinstein (1998). "Motives and Implementation: On the Design of Mechanisms to Elicit Opinions." *Journal of Economic Theory*, 79, 157–173.
- Kantor, Jodi (2006). "Entrees Reach \$40, and, Sorry, the Sides Are Extra." *The New York Times*, 21 October.
- Krishna, Vijay (2002). *Auction Theory*. Academic Press.
- Matsushima, Hitoshi (2007). "Mechanism Design with Side Payments: Individual Rationality and Iterative Dominance." *Journal of Economic Theory*, 133, 1–30.
- Matsushima, Hitoshi (2008). "Detail-free Mechanism Design in Twice Iterative Dominance: Large Economies." *Journal of Economic Theory*, 141, 134–151.
- McKelvey, Richard, and Thomas Palfrey (1995). "Quantal Response Equilibria for Normal-Form Games." *Games and Economic Behavior*, 10, 6–38.
- McMillan, John (1994). "Selling Spectrum Rights." *Journal of Economic Perspectives*, 8, 145–162.
- Myerson, Roger B. (1981). "Optimal Auction Design." *Mathematics of Operations Research*, 6, 58–73.
- Neeman, Zvika (2003). "The Effectiveness of English Auctions." *Games and Economic Behavior*, 43, 214–238.
- Pauzner, Ady (2000). "Independent Mistakes in Large Games." *International Journal of Game Theory*, 29, 189–209.