

Going. . . Going. . . Wrong: A Test of the Level- k (and Cognitive-Hierarchy) Models of Bidding Behavior

Itzhak Rasooly

Sciences Po and Paris School of Economics

In this paper, we design and implement an experiment aimed at testing the level- k model of auctions. We begin by identifying (simple) environments that optimally disentangle the predictions of the level- k model from the natural benchmark of Bayes-Nash equilibrium. We then implement these environments within a virtual laboratory in order to see which theory can best explain observed bidding behavior. Overall, our findings suggest that, despite its notable success in predicting behavior in other strategic settings, the level- k model (and its close cousin, cognitive hierarchy) cannot explain behavior in auctions.

I. Introduction

Although the study of auctions remains dominated by equilibrium-based approaches, in recent years more “behavioral” alternatives have emerged.

For useful comments and discussions, I would like to thank Johannes Abeler, Jesper Åkesson, Elizabeth Baldwin, Miguel Ballester, Adam Brzezinski, Riccardo Camboni, Vince Crawford, Jasmin Droege, Peter Eso, Caspar Jacobs, Bernhard Kasberger, Paul Klemperer, Nat Levine, Meg Meyer, Luke Milsom, Inés Moreno de Barreda, Frieder Neunhoeffler, Rossa O’Keeffe-O’Donovan, Karl Schlag, Alex Teytelboym, Séverine Toussaert, and Jasmine Theilgaard. I would also like to thank audiences at the Transatlantic Theory Workshop, SasCa PhD Conference, Oligo Workshop, and various seminars and workshops in Oxford. Finally, I am grateful to Darija Halatova and Tommaso Batistoni for coding assistance and to Exeter College and the Oxford Economic Papers Association for their generous financial support. This paper was edited by Eduardo Azevedo.

Electronically published May 12, 2023

Journal of Political Economy Microeconomics, volume 1, number 2, May 2023.

© 2023 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/723716>

Most prominent among these is Crawford and Iriberry (2007)'s level- k model, which posits that individuals iteratively best respond to a naive "first thought" about how one might bid in an auction. If accurate, the level- k model would substantially revise our understanding of how auctions ought to be designed (Crawford et al. 2009; De Clippel, Saran, and Serrano 2019) and would similarly revise how we think about individual reasoning in incomplete information settings. Understanding whether the level- k model is accurate is thus a critical issue.

Surprisingly, however, the existing literature provides little insight into the predictive performance of level k and whether it can outperform the natural benchmark of Bayes-Nash equilibrium (BNE). There is some evidence that the level- k model might be able to explain the "winner's curse" in common value auctions (Crawford and Iriberry 2007; Costa-Gomes and Shimoji 2015) but also evidence that it might do so for the wrong reasons (Ivanov, Levin, and Niederle 2010). In addition, there does not appear to be any serious comparison of level k and equilibrium in the rather simpler independent private setting (IPV) that forms the starting point for most auction research. For example, while Crawford and Iriberry (2007) do compare level k and equilibrium using an IPV data set, the auction structure that generated this data set leads level k and equilibrium to make nearly identical predictions—making a formal comparison of the models almost impossible.¹

It is also difficult to learn very much about the relative performance of level k and equilibrium using existing experimental data sets since these have often been generated by first-price auctions with uniformly distributed valuations (Kagel and Levin 1995). Unfortunately, the predictions of level k and equilibrium entirely coincide in such settings (for any $k \geq 1$). Moreover, while it is easy to disentangle level k and equilibrium using "exotic" value distributions, it is doubtful that such distributions would be understood by all subjects in any actual experiment. The challenge is therefore to find an experimental design that is both sufficiently simple for subjects to understand and yet dramatically disentangles the models.

¹ See also Kirchkamp and Reiß (2011) for a brief experimental comparison of level k and equilibrium, albeit in a paper not devoted to level- k reasoning. Since the goal of their paper was not to assess the performance of level k , it is perhaps not surprising that their experiments do not cleanly disentangle level k and equilibrium; e.g., in their first experiment, the predictions of the models coincide entirely. See also Gillen (2009), An (2017), Galavotti, Moretti, and Valbonesi (2018), and Hortaçsu et al. (2019) for work on the level- k model in the context of observational data. These papers do not attempt to test the accuracy of the level- k model or compare it with equilibrium, and indeed doing so using observational data (when individual valuations are unknown) would obviously be much more challenging than doing so with an experimental approach.

We begin by proposing an environment that does exactly this: a discrete all-pay auction with uniformly distributed values.² We first analyze the equilibrium predictions in this environment, proving that it possesses exactly one symmetric equilibrium. We thus extend a well-known result from continuous auction theory to the discrete auction model, and do so using completely novel (and elementary) arguments.³ In the course of our proof, we construct an algorithm that is guaranteed to identify this equilibrium,⁴ finding that it closely approximates the continuous equilibrium of traditional auction theory given the parameters chosen in our experiment. In this equilibrium, bids increase (quadratically) in valuations, rising to around half of valuations at the maximum possible valuation that players can draw.

We then conduct a similar analysis of the level- k model but uncover a surprising contrast. Solving the level-1 player's optimization problem reveals a corner solution: such types bid zero for all valuations. As a result, level-2 and level-3 types also bid close to zero; the former "overcutting" the level 1s by (almost always) bidding 1 and the latter overcutting the level 2s by (almost always) bidding 2. Clearly, this is very different from equilibrium, and so we have found an environment that clearly separates the models. In addition, this environment leads the models to possess very different comparative statics. For example, in the level- k model, making the bid discretization more coarse substantially inflates the predicted bids, whereas in the equilibrium model, it has very little effect.

Next, we conduct a similar exercise in the first-price auction, again assuming uniformly distributed values to make the environment experimentally implementable (and to force the two level-0 specifications proposed by Crawford and Iriberry [2007] to coincide; we elaborate on this point later on). To break the equivalence of the level- k and equilibrium models, we introduce the possibility that subject bids are "canceled" (in which case they lose the auction automatically). We then solve for the cancellation probability that maximizes the "distance" between the two theories. This turns out to separate the models in a similar way: level- k bids are once again close to zero, whereas equilibrium bids are roughly a quarter of valuations. Thus, we once again manage to disentangle

² We work with discrete models throughout both for the sake of realism (any actual auction must be discrete) but also because the level- k model would be ill defined in a continuous auction setting (we elaborate more on these points later on).

³ Our argument holds for arbitrary value distributions and can be straightforwardly modified to allow for factors such as risk aversion. It also extends immediately to the first-price auction; indeed, we suspect that it could be extended to fairly arbitrary auction structures (although we have not formally verified this).

⁴ To aid future experimenters, we have implemented the algorithm and made it available online. See <https://auctionsolver.herokuapp.com/>.

the models in an environment simple enough to be implemented in an experiment.⁵

We then conduct an experiment designed to put these insights into practice. The basic idea of the experiment is straightforward. We create the aforementioned environments in a virtual laboratory, allowing us to see which of the models under consideration better predict the observed bids. However, the experiment also possesses various other features aimed at testing the predictive success of the level- k model. For example, we elicit individuals' levels using a variant on Arad and Rubinstein (2012)'s 11–20 game, study whether behavior varies with the bid discretization in the manner predicted by level- k theory, and finally ask individuals to explain why they bid in the way that they did to see whether the level- k model accurately characterizes their conscious reasoning processes.

The experiment yields four central findings. First, if individual levels are restricted to the 1–3 range (to capture an intuitive limit on the number of thinking steps that individuals can be expected to perform), then the level- k predictions are far lower than the observed bids. In addition, the level- k predictions are substantially less accurate than those of equilibrium. These conclusions hold across all the auction structures and treatments that we consider and under a variety of approaches to model testing and comparison. Moreover, this conclusion continues to hold if one allows individual levels to vary across auction rounds but retains the requirement that individual levels are “plausible” (i.e., in the 1–3 range).

Second, if one fits the model without any such plausibility restrictions, one finds levels typically estimated in the 30–35 range. We view this almost as a *reductio ad absurdum* of the model since it conflicts with almost all prior experimental work on this issue (Crawford, Costa-Gomes, and Iriberri 2013), clashes with basic intuition about the number of thinking steps that individuals can plausibly be expected to conduct, and sits uneasily beside findings from linguistics that individuals struggle to even comprehend higher order statements whose order exceeds 3 (Arad and Rubinstein 2012). Moreover, methods for penalizing model flexibility (e.g., the Bayesian information criterion) generally do not support the inclusion of these very high levels.

Third, individuals' levels as estimated from the auctions bear almost no resemblance to their levels as inferred from a game known to trigger level- k reasoning (the 11–20 game). This is hard to square within the

⁵ One might wonder why level k and equilibrium can produce such different predictions in these environments when they produce identical predictions in the case of first-price auctions with uniform values. The answer is that, in the first-price auction with uniform values, the assumed level-0 behavior (uniform bidding) matches up with the equilibrium bid distribution (which is also uniform, albeit with a stretched support). In contrast, our two auction structures can lead to highly nonuniform equilibrium bid distributions—whereas the assumed level-0 behavior remains uniform.

context of the level- k model: for if levels reflect either cognitive sophistication or beliefs about the cognitive sophistication of one's opponents, then one would expect to find a positive correlation between estimated levels across games. Conversely, the absence of positive correlation across any two games can be taken as *prima facie* evidence that level- k reasoning is not operative in at least one of the games under consideration.

Finally, we find that subjects are unlikely to cite iterated reasoning when asked to explain why they bid in the way that they did. In other words, subjects are unlikely to offer conjectures about the bidding strategies of their opponents—and even less likely to ground these conjectures in what they take to be their opponents' conjectures about their own bidding strategy. Admittedly, we do see evidence of iterated reasoning in the explanations given by at least one subject—and this subject was (perhaps not coincidentally) likely to bid in line with the predictions of the level- k model. However, such reports are very rare, suggesting that the level- k model captures at best a small fraction of observed bidding behavior.

We then conduct a series of extensions and robustness checks. For example, we consider the impact of incorporating risk aversion into the models, allowing subjects to best respond to a distribution over levels (as in Camerer, Ho, and Chong [2004]), changing the level-0 specification, as well as conducting various tweaks to our empirical analyses (e.g., including subjects who submit dominated bids). We do not find that any of these exercises alter our substantive conclusions.

The remainder of this article is structured as follows. Section II outlines an equilibrium and level- k analysis of the discrete all-pay and first-price auctions. Section III presents our experimental design. Section IV contains our empirical results and section V our robustness checks. Finally, section VI concludes with a discussion of why the level- k model does so poorly in the auction setting despite its well-documented successes in other domains of strategic behavior.

II. Separating Level k from Equilibrium

A. *The All-Pay Auction*

To begin, let us describe our first environment that sharply disentangles the level- k and equilibrium models: a (discrete) all-pay auction with uniformly and independently distributed values. There are $n \geq 2$ bidders, each assumed to be risk neutral (we relax this assumption in sec. V). As in our experiment, each player's valuation v_i is drawn from the set $\mathbb{X} = \{0, 1, \dots, x\}$ for some $x \in \mathbb{N}^+$, and the bids that they may submit are restricted to the very same set \mathbb{X} . Let $\beta: \mathbb{X} \rightarrow \mathbb{X}$ denote a generic pure strategy (i.e., "bidding function"). We study the case of discrete values and bids since (i) this is inevitably the setup in the experiment that

follows, as indeed it must be in any auction experiment; (ii) if we had instead assumed that the bids were continuous, as is more standard in the literature, then the predictions of the level- k model would be ill defined.⁶

Since this is an all-pay auction, individuals need to pay their bid no matter whether they win the object. However, they win the object if and only if their bid is strictly higher than all bids submitted by their opponents. Note that we assume that individuals cannot win the object in the event of a tie. While this is a slightly unusual assumption, it matches the tie-breaking rule that we adopt in the experiment. This tie-breaking rule, in turn, was chosen because it tremendously simplifies the equilibrium (and to a lesser extent the level- k) analysis.⁷

To begin with, we consider what Bayes-Nash equilibrium (henceforth, “equilibrium”) predicts that individuals will do in this situation. In particular, we will focus on symmetric equilibria. We do this for three reasons. First of all, it is unclear how individuals could be expected to coordinate on asymmetric equilibria without any communication in a one-shot game. Second, symmetric equilibria are analytically more tractable. Third, based on computational experiments reported in Rasooly and Gavidia-Calderon (2020), it does not seem as though discrete auction games such as those described above possess asymmetric equilibria that are substantially different from their symmetric equilibrium. Therefore, one might hope that the restriction to symmetric equilibria is “almost” without loss of generality.

Unfortunately, since our auction is discrete we cannot directly apply results from “standard” auction theory. We can, however, establish the following very useful result (which holds entirely independently of the distribution from which each player’s valuation is drawn).

PROPOSITION 1. The discrete all-pay auction has exactly one symmetric equilibrium.

Proposition 1 contains two claims: that a symmetric equilibrium exists and that it is unique. Establishing the existence part is straightforward: since the game is finite, one can essentially just apply Harsanyi (1967). However, establishing uniqueness is considerably more involved; see appendix A for the details. Importantly, the uniqueness proof is based on an algorithm that explicitly constructs the (candidate) equilibrium (see app. B for examples of the algorithm in action). Thus, our result does

⁶ In particular, the level-2 player would not possess a best response: for any number $\epsilon > 0$ that they might bid, there exists a feasible bid $\epsilon' < \epsilon$ that would yield them a higher expected payoff.

⁷ While the tie-breaking rule assumed here mirrors that used in our experiment, one can obtain quantitatively similar equilibrium and level- k predictions using alternative tie-breaking rules (assuming that the number of bids and valuations is sufficiently large to make ties somewhat improbable).

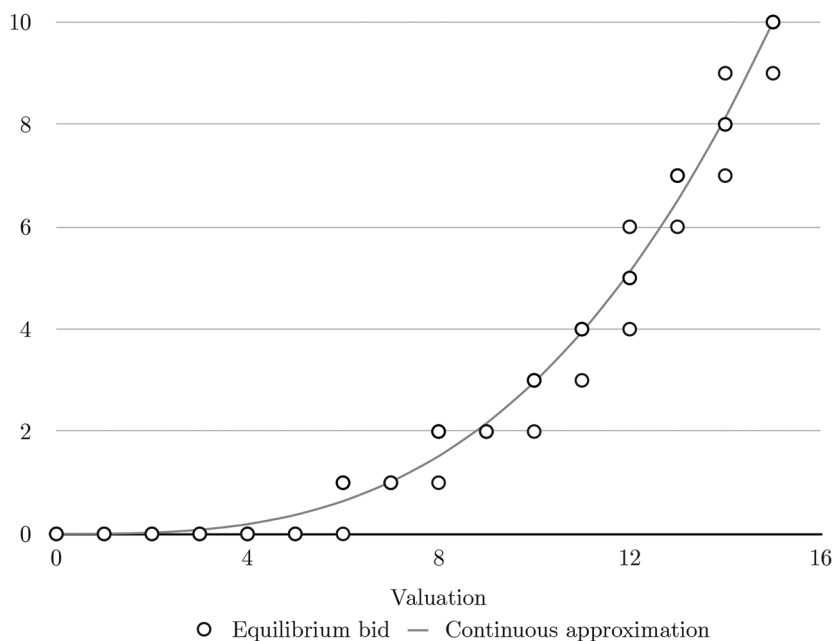


FIG. 1.—Equilibrium in the all-pay auction.

not merely tell us that there is exactly one symmetric equilibrium, but also shows us exactly how to compute it.⁸

Having developed this algorithm, we now use it to compute the symmetric equilibrium in the case of the distribution (uniform) used in our experiment. Figure 1 displays the result for the case of $x = 15$ and $n = 3$: a dot represents a bid that is submitted with positive probability. As can be seen, the equilibrium is in (properly) mixed strategies, but it is closely approximated, quantitatively speaking, by the continuous equilibrium of “textbook” theory:⁹

$$\beta(v) = \left(\frac{n-1}{n} \right) \frac{v^n}{x^{n-1}}. \quad (1)$$

As in the continuous model, equilibrium bids increase slowly when valuations are low but quickly when valuations are high. When values equal

⁸ This is invaluable from an experimental perspective since “brute force” approaches to equilibrium computation take several hours even for auctions with a handful of possible values and are essentially intractable for auctions of the size considered in our experiment (Rasooly and Gavidia-Calderon 2020). In contrast, our algorithm is able to compute the symmetric equilibria of our experimental games in less than a second.

⁹ See, e.g., Klemperer (1999) for a derivation.

the maximum x , bids are roughly equal to the equilibrium bid in the first-price auction; that is, $x(n - 1)/n$.

Having discussed the predictions of equilibrium (our main benchmark), we now turn to the predictions of the level- k model. As formulated by Crawford and Iriberri (2007), the model begins with a “naive” type of player (the level 0) who either bids their valuation (the “truthful” specification) or otherwise bids any element of the strategy space \mathbb{X} with equal probability (the “random” specification). This is meant to capture a player’s “first thought” about how their opponents will play the game. Higher types are then defined inductively: for $k \geq 1$, a level- k player submits the bid that maximizes their expected payoff given that their opponents are all level $k - 1$ (see Crawford and Iriberri [2007] for interpretation and elaboration). In the following, we will write $\beta^k : \mathbb{X} \rightarrow \mathbb{X}$ to denote a level- k player’s bidding strategy.

We now derive the behavior of the types. Assume, as in our experiment, that values are uniformly distributed on \mathbb{X} (retreating from the generality of proposition 1). Given the “truthful” specification ($b = v$), level-0 bids are then uniform on \mathbb{X} . Note, however, that this is also the “random” specification. We thus see that, quite aside from making the experimental rules easier for subjects to grasp, imposing uniform values ensures that the two level-0 specifications lead to identical predictions, a great advantage when it comes to testing the model.

Given that level-0 bids are uniform, and players win the auction if and only if all opponent bids are strictly below theirs, the level-1 player solves the problem

$$\max_{b \in \mathbb{X}} \pi(v, b) \equiv v \mathbb{P}(\text{win}|b) - b = v \left(\frac{b}{x + 1} \right)^{n-1} - b. \quad (2)$$

This leads to a corner solution: bidding 0 for all values $v \in \mathbb{X}$. To see why this is true, note that the benefit to such a player of increasing their bid by 1 is the gain in probability $\Delta \mathbb{P}$ multiplied by their valuation v . If $n = 2$, then $\Delta \mathbb{P} = 1/(x + 1)$, so the benefit of bidding one more is $v/(x + 1)$. Meanwhile, the cost of bidding one more is simply 1. Since $v/(x + 1) < 1$ for all $v \in \mathbb{X}$, the benefit of increasing one’s bid is always lower than the cost—which means that the optimal bid is 0. Moreover, if it is optimal to bid 0 when $n = 2$, it must be optimal for any $n \geq 2$: adding more bidders decreases the expected payoff from submitting any positive bid, while leaving the payoff from bidding 0 fixed.

It is then straightforward to derive the strategies of higher levels. For now, assume that a level- k player breaks ties by choosing the lowest optimal bid (while this assumption is natural since it allows us to avoid dominated bids, we relax it in sec. V). Given this assumption, a level-2 player sets $\beta(v) = 1$ for every value $v \geq 2$, and sets $\beta(v) = 0$ otherwise (they are

indifferent between bids of 0 and 2 when $v = 2$, so our tie-breaking rule comes into play). Similarly, a level-3 player sets $\beta(v) = 2$ for all $v \geq 3$, and sets $\beta(v) = 0$ otherwise. Thus, for all “reasonably low” levels, a level- k type player bids

$$\beta_{LK}(v) = \begin{cases} k - 1 & \text{if } v \geq k, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The strategy β_{LK} describes level- k behavior at “reasonably low” levels. At higher levels, however, this simple pattern breaks down. The strategy β_{LK} calls for a level- k player with a value of $k - 1$ to bid 0, thereby earning a (certain) payoff of 0. There is another option, however: bid 1 in the hope of winning when one’s opponents all bid 0. As k rises, this second option becomes more attractive, and eventually it pays to deviate from β_{LK} . It is easy to check, however, that level- k players must bid β_{LK} provided that

$$\pi(k - 1, 1) = (k - 1) \left(\frac{k - 1}{x + 1} \right)^{n-1} - 1 \leq 0 \iff k \leq (x + 1)^{(n-1)/n} + 1, \quad (4)$$

an inequality that gives concrete meaning to the phrase “reasonably low” in the previous paragraph.

Finally, we argue that level k can never coincide with equilibrium, with the implication that level- k bidding functions must cycle as $k \rightarrow \infty$. To prove this, note that by proposition 4 of Rasooly and Gavidia-Calderon (2020), the discrete all-pay auction with uniform values does not possess a symmetric equilibrium in pure strategies. Of course, it does have a symmetric equilibrium (see our proposition 1), so this must be in mixed strategies. Now, given our assumption about level- k tie breaking, it is clear that a level- k player could never choose a mixed strategy. As a result, for any $k \in \mathbb{N}^+$, level k can never coincide with the symmetric equilibrium. Moreover, this means that the level- k predictions must cycle. To see this, suppose that they did not cycle. Since the set of bidding functions is finite, there would need to exist some $k \in \mathbb{N}$ such that β^k coincides with β^k for all $k' > k$. (Otherwise, a strategy would need to reappear at some stage—which would necessarily trigger a cycle.) But then β^k would be a symmetric, pure strategy equilibrium, contradicting Rasooly and Gavidia-Calderon (2020)’s proposition 4.

The following proposition summarizes the findings of the previous paragraphs.

PROPOSITION 2. For every level k that satisfies (4) and every value $v \in \mathbb{X}$, a level- k bidder sets $\beta^k(v) = \beta_{LK}(v)$. Moreover, there is no $k \in \mathbb{N}$ at which β^k coincides with equilibrium, with the implication that β^k cycles as $k \rightarrow \infty$.

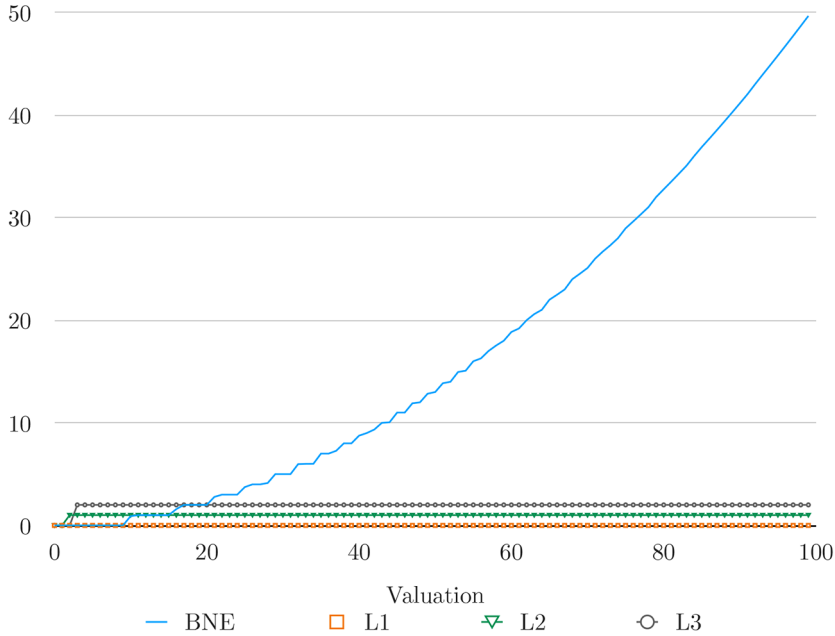


FIG. 2.—Comparing level k with equilibrium.

Figure 2 plots both the equilibrium predictions (or technically the continuous approximation) and level- k predictions (for levels 1–3). As can be seen, the auction structure has done an excellent job of disentangling the predictions of the models. First, given any reasonable calibration of the levels, the level- k model predicts bids that are substantially lower than those predicted by equilibrium. Second, while the equilibrium bidding function is generally increasing, the level- k bidding functions are flat at all but one value, providing a second contrast that one can test. Finally, our nonconvergence result means that the theories remain separated at arbitrarily high levels. It follows, then, that we have found a powerful way to disentangle the models, and to do so without abandoning the uniform-values specification that is so useful when implementing any actual experiment.

B. *The First-Price Auction*

It is clearly desirable to find an additional auction structure that sharply separates the models; that way, we can check the relative performance of the models in two different strategic settings. To this end, we now examine the possibility of doing this within the first-price auction. As noted by

Crawford and Iriberri (2007), the level- k and equilibrium models coincide (for every $k \geq 1$) in the first-price auction with uniformly distributed values. To break this equivalence, we now suppose that each player's bid only goes through with probability $p \in (0, 1)$; otherwise, their bid is canceled and they cannot win the auction. We now consider what probability p maximally disentangles the theories.

Formally, the model structure is much as before; for example, values and bids are restricted to the same set of consecutive integers \mathbb{X} . However, there are two differences. First, since this is now a first-price auction, a player pays their bid if and only if they win the object. Second, as just mentioned, each player's bid is now canceled with probability $1 - p \in (0, 1)$, and a player wins if and only if their (noncanceled) bid is strictly higher than all the noncanceled bids submitted by their opponents.

As before, we begin by demonstrating that a symmetric equilibrium exists and is unique:¹⁰

PROPOSITION 3. If $p \in (0, 1)$, then the discrete first-price auction has exactly one symmetric equilibrium.

The arguments used to establish this result are very similar to those used to prove proposition 1. In particular, the argument is based on the very same algorithm that constructs the equilibrium. As before, using this algorithm reveals that the equilibrium bids are well approximated by the continuous equilibrium, which is now (under uniform values)¹¹

$$\beta(v) = \left(\frac{n-1}{n}\right)v - \frac{x(1-p)}{np} \left[1 - \left(\frac{1-p}{1-p+p(v/x)}\right)^{n-1}\right]. \quad (5)$$

As can be seen, introducing the possibility that bids are canceled makes the bidding less aggressive. If $p = 1$, that is, bids are never canceled, our expression reduces to $\beta(v) = v(n-1)/n$. In other words, we recover the standard formula for equilibrium with uniformly distributed values. Moreover,

¹⁰ While our uniqueness result still holds under arbitrary value distributions, it now requires the restriction that $p < 1$; the role of this assumption is to ensure that all bids have positive probability of winning and thus to rule out equilibria involving dominated bids. Since this restriction holds in our experiment, this technical point need not concern us here.

¹¹ Equation (5) may be derived by noting the formal equivalence of our model with a model with an uncertain number of bidders. One can then apply theorem 1 of Harstad, Kagel, and Levin (1990). Alternately, one can note that our model is also equivalent (at least as far as undominated equilibria are concerned) to a model in which each player's value distribution has a mass of probability on zero. One can then use the "standard formula" (see, e.g., Krishna [2009])

$$\beta(v) = v - \int_0^v \frac{G(y)}{G(v)} dy,$$

and insert $G(v) = (1 - p + pv/x)^{n-1}$ (recalling that $G(v)$ is the probability that the maximum valuation is v or lower).

it is easy to check that for any $v \in \mathbb{X}$, $\beta(v)$ is strictly increasing in p . Finally, it can be shown (using L'Hôpital's rule) that $\beta(v) \rightarrow 0$ as $p \rightarrow 0$ for all $v \in \mathbb{X}$. This is also not surprising: if one is almost certain to win the auction anyway (in the event that one's own bid is not canceled), then one may as well bid $b \approx 0$.

We now turn to the predictions of the level- k model. Given that level-0 bids are either canceled or uniform on \mathbb{X} , a level-1 player solves the problem

$$\max_{b \in X} (v - b) \mathbb{P}(\text{win}|b) = (v - b) \left(1 - p + \frac{pb}{x + 1} \right)^{n-1}. \tag{6}$$

Since the objective function in (6) is single peaked, it is obvious that the integer that maximizes this function can be obtained by finding the $b \in [0, x]$ that maximizes the function and rounding up or down. In other words, we have a good continuous approximation of the exact level-1 bidding function. To simplify formulas, we will therefore proceed as if level-1 types could choose any $b \in [0, x]$ (although we will always use the exact model predictions when analyzing the data from our experiment). Routine optimization subject to the inequality constraint then reveals that

$$\beta^1(v) = \max \left\{ \left(\frac{n-1}{n} \right) v - \left(\frac{1-p}{p} \right) \frac{x}{n}, 0 \right\}. \tag{7}$$

As in equilibrium, level-1 bids are increasing in the probability p . Moreover, for all $p \leq 1/n$, the level-1 player sets $\beta(v) = 0$ for all $v \in \mathbb{X}$; that is, we recover the same corner solution as in the all-pay auction. Intuitively, if p is quite low, then there is a large chance that the player's opponents' bids are all canceled—so the player may as well just bid 0 and receive their entire valuation whenever they “get lucky.”

Our goal is to disentangle the models. To this end, we define the distance between the theories as

$$d(p) \equiv \int_0^x |\beta(v) - \beta^1(v)| dv. \tag{8}$$

Of course, the distance $d(p)$ is simply the total area between the (approximated) level-1 and equilibrium bidding functions. We now study the probability p^* that maximizes this distance. Recall from above that the level-1 player bids $\beta(v) = 0$ for all $v \in \mathbb{X}$ provided that $p \leq 1/n$. Meanwhile, every equilibrium bid $\beta(v)$ is strictly decreasing in p (moving closer to 0). From this, we see that we could never maximize (8) by choosing some $p \in [0, 1/n)$. The reason is that increasing p (by a sufficiently small $\epsilon > 0$) would increase every equilibrium bid $\beta(v)$ while leaving the level- k

bids unchanged, thus increasing the distance between the two theories. In other words, we have the following result.

PROPOSITION 4. $p^* \geq 1/n$.

Unfortunately, p^* cannot be solved analytically. Nonetheless, we can compute the optimal p numerically given various values of n (see app. C). We find, for example, that $p^* \approx 0.54$ if $n = 2$, $p^* \approx 0.34$ if $n = 3$, and $p^* \approx 0.26$ if $n = 4$. Overall, then, choosing $p = 1/n$ seems like a good approximation of the optimal probability (the loss in terms of distance is small). This has the further advantage of making the probabilities easier for subjects to grasp than the exact solutions would be, so we ultimately opt for $p = 1/n$ in our experiment.

Assuming that $p = 1/n$, it is then straightforward to derive the exact level- k bidding strategies. As noted, the level-1 player sets $\beta^1(v) = 0$ for all $v \in \mathbb{X}$. As a result, the level-2 player bids 1 (when their value is high) or bids 0 (when their value is low). Similarly, the level-3 player bids 2 for high values and 0 otherwise. This simple pattern characterizes the behavior of level- k bidders when k is “reasonably low” relative to the number of valuations $x + 1$. More formally, we can say the following:

PROPOSITION 5. Fix a set of types $\mathbb{K} = \{1, 2, \dots, \bar{k}\}$, let $p = 1/n$, and define

$$v^*(k) = \frac{k - 1}{1 - (1 - p)^{n-1}}.$$

Then if x is sufficiently large, every level- k type in \mathbb{K} bids $\beta^k(v) = k - 1$ for $v > v^*(k)$ and $\beta^k(v) = 0$ otherwise.

Given the parameters ultimately chosen for our experiment ($n = 2$, $p = 1/2$, $x = 100$), level- k behavior once again cycles at very high levels. To illustrate this, figure 3 plots the maximum bid submitted by the first 100 levels (the full level- k bidding functions are available in the supplementary materials). As can be seen, the model produces cycles of length 29 (so level 7 coincides with level 36, which in turn coincides with level 65, etc.). Moreover, far from cycling locally around the equilibrium predictions, level- k cycles are quite dramatic: the maximum bid ranges from 0 to 20. This is all the more surprising once we realize that all bids above 50 are strictly dominated (by bidding zero), so the highest level- k bid traces out around 40% of the nondominated strategy space.

Finally, figure 4 compares the models in the case of $n = 2$ and $p = 1/2$. As can be seen, we have once again managed to disentangle the theories—and in a dramatic way. First, given any plausible calibration of the levels—recall that these are supposed to be in the 1–3 region—level- k bids are significantly lower than equilibrium bids. Second, while equilibrium bids are increasing in values, level- k bids are almost always nonincreasing. Finally, the nonconvergence of level k to equilibrium means that the models

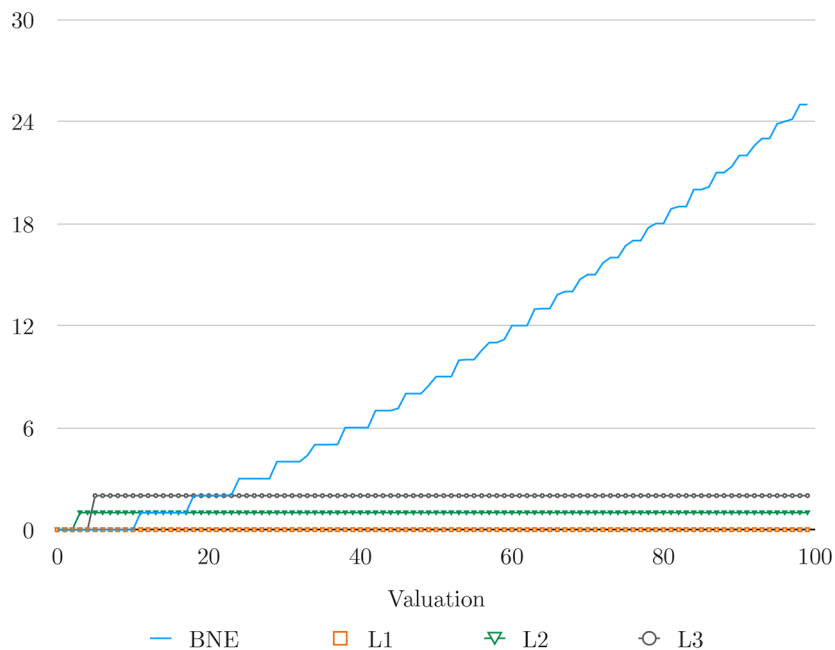


FIG. 4.—Comparing level k with equilibrium.

auction (as described in the previous section) with $n = 2$. Following the auctions, subjects played a modified version of Arad and Rubinstein (2012)'s 11–20 game that is proposed in Alaoui and Penta (2016). The purpose of this game was to allow for a rough calibration of each participant's level.¹³ Finally, we measured each subject's risk aversion using the “bomb” risk elicitation task (Crosetto and Filippin 2013).

Before participating in either of the auctions, subjects were given extensive quizzes to check their understanding of the rules. They were not permitted to proceed past the quiz until they had answered a full set of questions correctly. We view this as an important feature of the study: subjects are unlikely to make meaningful choices if they cannot even understand the rules of the game that they are playing.

Each subject bid in two rounds of each of the auction types.¹⁴ Moreover, we did not inform subjects of the bids submitted by their opponents in

¹³ We use Alaoui and Penta (2016)'s modification of the original game since it prevents cycling (level- k behavior coincides for all $k \geq 9$). As Alaoui and Penta (2016) note, this is then useful if one wishes to calibrate individual levels without the potentially controversial assumption that each individual's level is the smallest possible level that rationalizes their observed behavior.

¹⁴ They were also asked to explain the general considerations determining how they chose to bid after every second round.

previous rounds. We made these two choices—only two rounds coupled with very little feedback—since the level- k model is supposed to be a model of initial play, as opposed to a model of the outcomes of a long-run learning process (Crawford and Iriberry 2007).¹⁵ In other words, we sought to test the level- k model within its intended domain of applicability.

Subjects were informed (and regularly reminded) that they would never bid against the same opponent twice across the rounds of an auction. That is, we used the “perfect stranger matching” protocol. We opted for this protocol since the two theories under consideration are static theories—in the sense that they treat the game as a one-shot interaction—and we did not want to complicate the analysis by introducing learning or reputational considerations.

Bidding was incentivized through a procedure similar to that employed in Filiz-Ozbay and Ozbay (2007). In every round, individuals were asked how much they would want to bid given 10 possible valuations that they might have. They were informed that they should bid carefully since one of these valuations would be their actual valuation and they would be committed to bidding the amount they had entered. In other words, 1 in 10 bids were selected to “count.” We opted for this procedure, as opposed to incentivizing every bid, since it allowed us to collect more data given a fixed number of rounds, thereby further reducing feedback or learning effects (and making the environment more hospitable to the level- k model).¹⁶

Finally, we randomly allocated the participants into two treatments: an “integer bid” treatment (to which 2/3 of participants are allocated) and a treatment in which bids were constrained to be multiples of 5. The idea here was to check whether varying the bid treatment has the strong effects predicted by level- k theory or the relatively weak effects predicted by the equilibrium model. In the level- k model, increasing the discretization by a factor of 5 increases the predicted bids by roughly a factor of 5, assuming that levels are sufficiently low for individual behavior to be characterized by propositions 2 and 5 (see the supplementary materials for an exact description of the effects of the bid discretization on the model’s predictions). In contrast, varying the bid discretization does not have large quantitative effects on equilibrium bids; for example, in the first-price auction the average equilibrium bid increases by just 4.7% (again, see the supplementary materials).

The experiment was preregistered along with a full analysis plan.¹⁷ The experiment also received ethics approval from both the University of

¹⁵ Indeed, it is for this reason that Crawford and Iriberry (2007) only test their model using data from the first five rounds of the experiments they examine.

¹⁶ On average, subjects earned £9.60 over the course of around 40 minutes. Total earnings ranged from £3.60 to £16.00 with a standard deviation of £2.34.

¹⁷ See <https://www.socialsciregistry.org/trials/8011>.

Oxford's and Centre for Experimental Social Sciences' institutional review boards.

IV. Results

A. Pooling the Data

We begin by asking which theory best explains the overall patterns observed in the bidding data. To this end, we pool all the valuations and bids, compute the mean observed bid for every valuation, and plot this against the corresponding predictions of the two models. For the moment, we assume that levels cannot exceed 3,¹⁸ and drop the handful of subjects who submitted at least one dominated bid in one of the auctions.

Figure 5 displays the results for the case of the first-price auction with integer bids (see also fig. D1 for the treatment in which bids are multiples of 5). As found in many prior studies of first-price auctions, we observe substantial overbidding relative to the (risk-neutral) equilibrium benchmark (Kagel and Levin 1995). Indeed, the average bid is around 46% of the average value, whereas equilibrium predicts that it should be around 20%. However, while the equilibrium model struggles to match the data, it is clear that the level- k model performs considerably worse, predicting bids that are at least an order of magnitude lower than those commonly observed. As a result, it is clear that equilibrium provides a substantially better (albeit highly imperfect) description of the average bidding function observed in the experimental data.

We then conduct a similar analysis of the all-pay auction. As shown in figure 6, equilibrium does a somewhat better job of tracking mean bids in this environment (see also fig. D2 for the case in which bids are multiples of 5). However, overbidding relative to equilibrium is again observed on average at low to medium values.¹⁹ As before, the level- k model predicts bids that are orders of magnitude below those that are observed on average, at least if one imposes the restriction that levels cannot exceed 3. Thus, our initial examination of the data suggests a very clear ranking of the two models under consideration.

To verify these conclusions more formally, we also compute the root-mean-square prediction errors of both of the theories. In the case of equilibrium, this is straightforward since the model is "parameter free." In the

¹⁸ We impose this constraint because all of the evidence of which we are aware suggests that insofar as individuals use level- k reasoning, it rarely goes beyond the third level. Indeed, this is what we find when examining the responses of our subjects in the 11–20 game: out of those who mentioned level- k reasoning when asked to explain their choices, not one made a choice consistent with more than three levels of reasoning.

¹⁹ The overbidding that we observe is consistent with results from the only other experimental study of IPV all-pay auctions of which we are aware (Noussair and Silver 2006). Unfortunately, however, we were unable to obtain the data from this experiment and so were unable to use it to complement our analyses.




FIG. 5.—Predictions and data in the first-price auction.

case of level k , we obtain a prediction by assigning each data point the level from the 1–3 range that minimizes the model’s prediction error (this is a generous assumption that will tend to overstate the model’s predictive performance). Table 1 displays the results. As can be seen, both models make sizeable prediction errors. Consistent with the graphical evidence, however, the prediction errors of level k are substantially greater than those of equilibrium.

Next, we examine the effect of varying the bid discretization. Recall that the level- k model predicts a large effect from changing the discretization from integers to multiples of 5: bids should increase by roughly a factor of 5. In contrast, the equilibrium model predicts much smaller effects. Given these very different predictions, we now conduct a comparison of the integer bid and multiples of 5 bid treatments.

As a preliminary, we confirm that the treatments are balanced, suggesting that the randomization worked as expected (see table D1). Next, we inspect the average bids across the two treatments and auction structures (see table D2). As can be seen, changing the bid discretization makes essentially no difference to the first-price auction (in both cases, the average bid is about 46% of the average valuation). In the all-pay auction, restricting bids to multiples of 5 does appear to slightly increase the average




FIG. 6.—Predictions and data in the all-pay auction.

bids (which rise from 45% to 50% of the average valuation), but clearly not in the dramatic way predicted by the level- k model. In neither auction, then, do the data support the surprising comparative static prediction of level- k theory.

B. Structural Estimates

In light of the evidence above, it is clear that the level- k model cannot account for the observed average behavior, at least when calibrated with a plausible distribution of types. However, this still leaves at least two issues unresolved. First, even if level k cannot describe the average bidder, it is still possible that it accounts for a substantial minority of the observed

TABLE 1
PREDICTION ERRORS

	T1 FP	T1 AP	T2 FP	T2 AP
Equilibrium	15.1	13.3	14.1	14.3
Level k	21.3	19.8	21.4	24.0

NOTE.—T1 and T2 denote the integer bid and multiples of 5 treatments; FP and AP abbreviate “first price” and “all pay.”

behavior. Second, it remains to be seen whether the model might be rescued by allowing for a greater number of possible types. We now address both of these issues.

To do this, we follow the literature in estimating a mixture of types model (see, e.g., Stahl and Wilson [1994], Bosch-Domènech et al. [2002], and Crawford and Iriberrri [2007]). To obtain stochastic versions of the relevant theories, we assume that subjects submit bids with probabilities that are proportional to a normal distribution whose mean is the predicted bid that they are “supposed” to submit (according to the relevant theory). More formally, let $\mathbb{K} \subset \mathbb{N}^+$ denote the set of types, indexing equilibrium by $k = 0$ and level- k types in the obvious way. Given that a player i has value $v_i \in \mathbb{X}$ and type $k \in \mathbb{K}$, we assume that the chance they submit a bid $b_i \in \mathbb{X}$ is given by

$$\mathbb{P}(b_i|v_i, k) \propto \exp\left[-\frac{1}{2}\left(\frac{b_i - b_p}{\sigma_k}\right)^2\right], \tag{9}$$

where b_p is the bid predicted by their type (given their valuation), σ_k is a noise parameter that controls the variance of their bids, and the constant of proportionality is chosen to ensure that the probabilities sum to 1.²⁰ This produces a single-peaked distribution whose mode is the predicted bid b_p . As $\sigma_k \rightarrow \infty$, this distribution converges to uniform noise, and as $\sigma_k \rightarrow 0$, it converges to a degenerate distribution that places all probability mass on the predicted bid b_p .

For now, assume that each individual’s type is fixed (though we relax this in many of our analyses). Then the probability that an individual submits a vector of 20 bids \mathbf{b}_i given their vector of values \mathbf{v}_i is (assuming independence of errors)

$$\mathbb{P}(\mathbf{b}_i|\mathbf{v}_i, k) = \prod_{i=1}^{20} \mathbb{P}(b_i|v_i, k). \tag{10}$$

We assume that the individual we observe is drawn randomly from the population, and write $p_k \in [0, 1]$ to denote the population proportion of type k (obviously, $\sum_k p_k = 1$). By the law of total probability, the unconditional probability of their bids is

$$\mathbb{P}(\mathbf{b}_i|\mathbf{v}_i) = \sum_{k \in \mathbb{K}} \mathbb{P}(\mathbf{b}_i|\mathbf{v}_i, k)p_k, \tag{11}$$

and so the “likelihood” of all the observed decisions is (combining eqq. [9]–[11])

²⁰ We also experimented with a model containing logit errors. However, these had the unfortunate tendency of preferring level 1 to all higher levels even though level-1 predictions were the farthest from the observed bids (recall that level-1 types bid zero regardless of their valuation).

$$L = \prod_{i=1}^n \mathbb{P}(b_i | v_i) \propto \prod_{i=1}^n \left(\sum_{k \in \mathbb{K}} \prod_{j=1}^{20} \left(\exp \left[-\frac{1}{2} \left(\frac{b_i - b_p}{\sigma_k} \right)^2 \right] \right) \right) p_k. \quad (12)$$

In every estimation exercise, we compute the parameters $p = (p_k)_{k \in \mathbb{K}}$ and $\sigma = (\sigma_k)_{k \in \mathbb{K}}$ that maximize the likelihood (12). We then calculate standard errors with the jackknife. While we consider the case in which individual levels are fixed (as above), we also consider the case in which individual levels are permitted to vary by round—although we do insist throughout that levels are fixed within a round.²¹

As a preliminary exercise, we compare an equilibrium-only model with the level- k model in which levels are restricted to 1–3. To avoid introducing too great an imbalance in the number of parameters between the two respective models, we restrict both to a single noise parameter σ . Table D3 displays the results. As one might have expected from figures 5–6, equilibrium does a better job than any of the level- k models at explaining the data, attaining a higher log likelihood in all specifications. This advantage becomes even greater if we punish the level- k models for their extra flexibility, for example, by using the Bayesian information criterion. This simply reinforces the conclusion that we reached earlier: equilibrium does a better job than the level- k model at accounting for the overall patterns in the data.

We now examine the extent to which the level- k model might be salvaged by adding higher levels. To this end, we now estimate a series of models, gradually adding levels until we reach the point where the model cycles. Three main points emerge from this exercise.²² First, while the level- k model can achieve a higher likelihood than equilibrium once the full suite of levels is added, it can do this only by placing most probability on very high levels; for example, in the first-price auction, most probability mass is placed on level 32. As argued previously, is doubtful whether we can plausibly ascribe so many thinking steps to participants. Second, the level- k model fits the data with a very unusual distribution, placing some probability mass on level 1, none on intermediate levels, and the rest on the very high levels just discussed. Of course, the idea that individuals do either one thinking step or 32 thinking steps might appear to be rather strange. Finally, the inclusion of these very high levels is almost never recommended by the Bayes information criterion. For all these reasons, it seems difficult to rescue the level- k model through a strategy of adding a very large number of levels.

²¹ Recall that all bids within a round must be made simultaneously, so there is no possibility for learning within rounds. In contrast, individuals could conceivably learn between rounds—although the scope for such learning is limited by the near-total absence of feedback.

²² Since the resulting tables are too large to be contained within the appendix, they have been made available at https://docs.google.com/spreadsheets/d/1qNrerSNlzhqOSkyISZL29tgxcjxqW_3wO3refseew3o/edit.

Although a pure level- k model appears to be outperformed by equilibrium, a final question we investigate is whether a hybrid model might be able to outperform both. To this end, we now reestimate the model for all auction structures, allowing for both equilibrium and level- k types and no longer restricting the model to a single noise parameter σ . We also return to our previous assumption that levels are in the plausible range (1–3). Table D4 displays the results. As can be seen, most probability mass is placed on equilibrium—suggesting again that the level- k model fits the data relatively poorly. However, there is still some probability mass placed on level- k types. Thus, the results from the hybrid model are consistent with the possibility that the level- k model characterizes the behavior of a small minority of subjects, even if it struggles to capture the general patterns in the bidding data.

C. *Correlating Levels*

We now examine whether individual levels as estimated from the two auctions bear any resemblance to the levels as inferred from a game known to trigger level- k reasoning: a variant on the 11–20 game proposed by Alaoui and Penta (2016). To this end, we calculate the level k and precision σ_k that maximizes the likelihood of each individual's bids (i.e., eq. [10]). Thus, we now assign each individual a definite level, as opposed to a distribution over levels, and no longer allow for equilibrium types. After doing this for both the first-price and all-pay auction, we then compute each individual's level as inferred from the 11–20 game, discarding those whose implied level exceeds 4.²³ Finally, we compute the correlation between the levels inferred from each auction and the level as inferred from the 11–20 game.

Figure D3 plots the levels estimated from the first-price and all-pay auctions against the levels as inferred from the 11–20 game. As can be seen, there is no discernible relationship between the levels, and indeed the correlations are estimated to be 0.00 and -0.09 in the case of the first-price and all-pay auctions, respectively.²⁴ From the perspective of level- k theory, this is rather unexpected. For if levels reflect either cognitive sophistication or beliefs about the cognitive sophistication of one's opponents—the two leading interpretations of individual levels—then one would expect to find a positive correlation between individual levels across games. Conversely, the absence of such a positive correlation

²³ We do this partly because such levels are at odds with prior experimental evidence but also because not one of such subjects invoked level- k reasoning when asked to explain their choice.

²⁴ This result reflects those of Georganas, Healy, and Weber (2015), who find no correlation between estimated individual levels across two families of games (albeit in a non-auction setting).

suggests that individuals are not using level- k reasoning in at least one of the games.²⁵

D. Subject Reports

Finally, we examine the explanations given by participants as to why they chose to bid in the way that they did. In particular, we consider whether they appealed to iterated reasoning.²⁶ Assuming that subjects' conscious reasoning was indeed captured by the level- k model and further assuming that subjects did not for some reason wish to lie to their experimenter about their reasoning, one would expect to see iterated reasoning cited in their explanations. Conversely, one can take the absence of such reports as evidence that the level- k model did not in fact characterize their conscious reasoning.

In their explanations, the majority of subjects (around 55%) explicitly indicated an awareness of the basic bidding trade-off: bidding more increases one's probability of winning, but decreases one's payoff if one wins the auction. For example, one subject reported trading off "the chance of winning versus the reward for winning." However, the vast majority of subjects did not appear to use iterated reasoning.²⁷ Indeed, we categorized responses as invoking iterated reasoning in just two cases.²⁸

To be clear, we are not claiming that the level- k model did not characterize the reasoning given by any of the subjects. For example, one subject (who also bid in line with the model) wrote as follows:

²⁵ One can also ask whether subjects who made different choices in the 11–20 game made different choices in the auctions (whether or not these choices match up with the predictions of level- k theory). As shown by table D5, subjects who chose 17 (level 3) in the 11–20 game tended to bid more aggressively in both auctions than the rest, although there is little difference in average bids between subjects who chose 18 (level 2) and those who chose 19 (level 1).

²⁶ A strict definition of iterated reasoning would be reports of the following form: "I bid in way X because I expected my opponent to bid in way Y, and I expected my opponent to bid in way Y because I thought they would think I would bid in way Z." In other words, this would need to involve (1) a conjecture about the bidding strategies followed by the player's opponents and (2) an attempt to ground this conjecture using a conjecture about the conjectures of the opponents about the player's bidding strategy. On a weaker definition, we would simply require some kind of conjecture about the opponents' bidding strategy (along with an attempt to optimize given this conjecture). Since the level- k model allows for level-1 players, we used this weaker (and more permissive) definition when categorizing the reports.

²⁷ All categorizations were independently checked by another researcher who was not involved in the study. The reader is invited to examine the responses themselves, which are provided (along with all other nondemographic data generated by the experiment) in the supplementary materials.

²⁸ In fairness, one should point out that subjects also did not invoke the kind of fixed-point reasoning presupposed by equilibrium. Thus, it seems doubtful that either equilibrium or level k adequately characterizes the conscious reasoning process of most subjects.

Since my opponent has a 50% chance of their bid being cancelled, it's better to keep bids low in hopes theirs gets cancelled. . . . I raised the bid to more than 1 because in the event my opponent is using the same strategy and bids very low, I still have a chance of winning the auction.

Thus, this subject not only considered something like the level-1 strategy (bidding “very low”) but then went beyond it—a clear case of level- k reasoning.²⁹ What is evident, however, is that reasoning was provided by only a small minority of subjects, a finding that is also in line with the choice data outlined in the previous section.

V. Robustness Checks

We now report on the results of a variety of extensions and robustness checks.

Risk aversion.—While the previous analysis assumed that individuals are risk neutral, we now consider whether the level- k model might be salvaged by allowing individuals to be risk averse.³⁰ To begin, consider the all-pay auction. As noted in section II.A, level-1 players (if assumed to be risk neutral) are predicted to bid 0 (for all values), thereby earning a certain payoff of zero. If they submitted a positive bid, then they would be exposed to risk: either their payoff would be positive (if they win the auction) or negative (if they lose the auction). Now, given that level-1 players choose to bid 0 when they are risk neutral, it is obvious that they must also choose to bid 0 if they are assumed to be risk averse. Thus, introducing risk aversion makes no difference to the model's predictions—and also (for this reason) makes no difference to the predictions for high levels (provided that they are sufficiently low to be characterized by bidding strategy β_{LK}).

While risk aversion cannot improve the model's performance in the all-pay auction, matters are more complicated in the first-price auction. In general, holding the bidding strategies of one's opponents fixed, introducing risk aversion must (weakly) increase one's optimal bid. In our setting, this means that introducing risk aversion can make positive bidding by a level-1 type optimal, thereby bringing the model's predictions closer in line with the observed data. Note, however, that introducing risk aversion also makes equilibrium bids more aggressive (Holt 1980), which

²⁹ While the bids submitted by the subjects did not exactly conform to any of the levels, they were very close, with the subjects bidding a mix of 2's and 3's at relatively low valuations.

³⁰ Of course, expected utility maximizers should be approximately risk neutral over such small stakes (Arrow 1970; Rabin 2000). Thus, the “risk aversion” discussed here should really be thought of as capturing utility curvature within a reference-dependent model (Bleichrodt et al. 2019).

might be expected to increase the predicted accuracy of our benchmark. It is therefore unclear at the onset whether introducing risk aversion improves the relative performance of the level- k model.

To study this issue more carefully, we assume that every individual i has constant relative risk aversion (i.e., $u_i(x) = x_i^\alpha$, where x_i is the individual's payoff and $\alpha_i > 0$ determines their level of risk aversion). We then estimate every individual's parameter α_i using their choice in the "bomb" risk elicitation task (Crosetto and Filippin 2013). We then recalculate the level- k model's predictions for every player using their estimated parameter α_i (along with their estimated level as in sec. IV.A). We also recalculate the equilibrium predictions on the assumption that every individual has the average estimated risk aversion parameter of the subject pool ($\alpha = 0.72$).³¹

Figure 7 displays the results for the main treatment (see also fig. D4 for the case in which bids are multiples of 5). As can be seen, both the level- k and equilibrium predictions are now somewhat improved and considerably closer to the actual data. Overall, however, it is clear that equilibrium still outperforms level k dramatically, a conclusion confirmed by the prediction error rates reported in table D6.

Cognitive hierarchy.—We now ask whether any conclusions might change if individuals are assumed to best respond to a distribution over levels, as in Camerer, Ho, and Chong (2004)'s cognitive-hierarchy model. To this end, we assume (following Camerer, Ho, and Chong [2004]) that levels are Poisson distributed with mean τ . Using standard maximum likelihood methods, we then estimate τ for the participants in our experiments using their choices in the 11–20 game, finding that $\tau \approx 1.9$. Finally, we compute the predictive errors of the cognitive-hierarchy model and compare those with the associated errors for equilibrium. As shown in table D6, moving from level k to cognitive hierarchy makes essentially no difference to the prediction errors. This is not a surprise: obviously, best responding to a mix of levels between 1 and 3 yields (in our context) very similar optimal bids to best responding to level 2.

Including dominated bids.—As noted earlier, our main analysis excluded the small number of subjects who submitted at least one dominated bid on the grounds that such subjects may have misunderstood the rules of the auction. However, excluding such subjects is a debatable choice, and we reran all analyses after including such subjects. This did not change any of our substantive conclusions (see table D6).

Changing the tie-breaking rule.—In the previous analysis, we assumed that level- k players always chose the lowest optimal bid if multiple bids were

³¹ The code can be viewed here: https://github.com/Itzhak95/risk_aversion/tree/master. The $\alpha = 0.72$ is obtained after dropping the two outliers who collected 61 or more out of the 62 boxes. If such subjects are included, we obtain $\alpha = 1.8$ (risk seeking), in which case allowing nonlinear utility is unable to improve the performance of the level- k model.




FIG. 7.—Risk aversion in the first-price auction.

optimal. While this allowed the model to avoid predicting that individuals choose dominated strategies, it is natural to ask whether any of our conclusions hinge on this assumption. To investigate this, we now make the opposite assumption—that level- k players always choose the highest bid out of the bids that they believe are optimal.³² Rerunning our main analyses, we see that this slightly worsens the model’s fit with the data (although the change is small). This is not a surprise: the main change is that the model now predicts some dominated bids, and such bids are very rare in the data.

Dropping the second round.—As noted in section III, the level- k model is supposed to be a model of initial play. In order to test the model in its natural domain, we therefore blunted subject feedback and restricted the number of rounds to two. Nonetheless, one might worry that even two rounds are too many to count as “initial play.” To test this possibility, we reran the analyses after dropping the second round. While this does somewhat improve level k ’s performance, the model remains clearly outperformed by equilibrium (see table D6).

³² Obviously, this does not exhaust the universe of possible tie-breaking rules: one could also assume that players sometimes choose the highest optimal bid and sometimes choose the lowest optimal bid.

Changing the level-0 specification.—The level- k auction model, as formulated by Crawford and Iriberri (2007), assumes that level-0 players either bid their value (the “truthful” specification) or randomize uniformly over \mathbb{X} . Since the predictions of such a model seem strongly disconfirmed by the data, one might wonder whether the model might be improved by choosing an alternative level-0 specification. While our data do not allow us to absolutely rule this possibility out, there are various reasons to think that this is not likely to be a fruitful theoretical exercise.

First, as stressed by Crawford, Costa-Gomes, and Iriberri (2013), one must be careful to ensure that any level-0 specification represents a “strategically naive” assessment of how others will play the game. This is most obviously important if one wishes to connect the model to its standard interpretation as a model of iterated reasoning given some naive but initially plausible starting point. In addition, it is vital to restrict the range of possible level-0 strategies if one wishes the model to make anything like definite predictions in strategic situations, and to be clearly disentangled from rival models. For instance, if an analyst were allowed to specify that level-0 play coincided with the symmetric equilibrium, it would be impossible to disentangle the equilibrium from the predictions of the level- k model: for every $k \geq 1$, the level- k prediction would simply be the symmetric-equilibrium strategy.

Thus, if one did want to change the level-0 specification, one would need to make sure that it represents a psychologically plausible way in which a player might naively expect others to play. This brings us to our second point, however: there simply do not appear to be a large number of such level-0 specifications. Other than the specifications proposed by Crawford and Iriberri (2007), we are only able to think of one alternative: bidding uniformly between 0 and one’s value v (modifying their “random” specification so as to remove dominated bids). However, opting for this level-0 specification would render the model unable to match even the general trend of bidding in the most experimentally studied of all auction formats—the first-price auction with independent, uniform values (and no possibility of canceled bids).³³

³³ In this setting, a large number of auction experiments report overbidding relative to the risk-neutral equilibrium $\beta(v) = (n-1)v/n$ (see Kagel and Levin [1995] for a review). However, it is easy to see the level-0 specification suggested here would predict bids that are substantially lower than the risk-neutral equilibrium. As a result, it would result in level- k predictions that are far too low (at least, given a plausible calibration of the levels). Note also that other apparently plausible level-0 specifications, such as thinking that one’s opponents bid half their value, are excluded on the grounds that level-0 reasoning is supposed to be nonstrategic (whereas bidding a fraction of one’s value quite clearly reflects an attempt to strategically balance one’s chance of winning with one’s payoff if one does win). Moreover, such specifications imply that level- k bidders should submit a large mass of bids just above the maximum value bid by level 0 (e.g., L_1 players in the all-pay auction should bid 50 for all $v \geq 51$). We do not observe such masses in this experiment or in the other auction experiments of which we are aware.

Finally, we note that the level- k predictions as formulated in section II are robust to a very wide range of level-0 specifications. For example, in the all-pay auction, level- k bids will be characterized by proposition 2 for any level-0 specification that satisfies

$$F(b) \leq \left(\frac{b}{x+1} \right)^{1/(n-1)}, \tag{13}$$

where $F(b)$ is the probability that the level-0 type bids b or lower (i.e., the cumulative distribution function [CDF] of the bid). In particular, this means that proposition 2 will continue to hold provided that the CDF $F(b)$ is bounded by the uniform CDF from above. Similarly, in the first-price auction, level-1 players will bid 0 (as in proposition 5) for any level-0 specification that satisfies

$$F(b) \leq \left(\frac{n-1}{n} \right) \left[\left(\frac{x}{x-b} \right)^{n-1} - 1 \right] \tag{14}$$

for all $b \in \mathbb{X}$.³⁴ Moreover, while these inequalities are sufficient for our level- k characterization to hold, they are absolutely not necessary. For example, in the all-pay auction, the level-1 bidding strategy turns out to be a best response to the empirical bid distribution observed in the experiment, even though that distribution does not satisfy inequality (13).

In light of the reasons above, we think that it is unlikely that one can find a level-0 specification that is simultaneously nonstrategic, psychologically plausible, and generates predictions that match the data in both this and other auction experiments. Ultimately though, this conclusion is somewhat more speculative than the central conclusion of this study: that the level- k model as specified by Crawford and Iriberry (2007) cannot predict or rationalize the observed bidding behavior.

VI. Concluding Remarks

In this paper, we have designed and implemented an experiment aimed at testing the level- k model of auctions. Overall, the evidence would appear to strongly reject the model. When plausibly calibrated, the model produces bids that are at least an order of magnitude too low. Moreover, fitting the model without such a plausibility constraint results in estimated levels in the 30–35 range, and these levels in turn fail to bear any relation to levels inferred from the 11–20 game. In addition, changing the bid discretization has little effect on the observed bids (in contrast to the model’s predictions), and subjects almost never cite iterated reasoning

³⁴ This inequality comes from checking that a player with $v = x$ would rather bid 0 than any $b \in \mathbb{X}$.

when asked to explain why they bid in the way they did. Thus, while one might be able to disagree with any of the pieces of evidence separately, it seems that the evidence collectively casts strong doubt on the model's ability to predict or even rationalize the observed data.

This finding might be seen as surprising given that the experimental environment was carefully chosen to be the level- k model's natural domain of applicability. Individual feedback was severely restricted and (following Crawford and Iriberri [2007]) the number of rounds was restricted to just two. However, the level- k model fared poorly despite the initial-play setting.

The findings also might be seen as surprising in light of evidence that the level- k model very accurately captures facets of human behavior in other settings, whether "beauty contests" (Nagel 1995), coordination games (Costa-Gomes, Crawford, and Iriberri 2009), strategic communication (Crawford 2003; Cai and Wang 2006) or zero-sum betting (Brocas et al. 2014). We thus close with some speculations as to why the model appears to do so badly at predicting bids in auctions despite its successes in predicting behavior in other areas.

First, recall that the level- k model requires a level-0 anchor, that is, some sort of naive starting point that forms the basis of the iterated reasoning. In many settings, such an anchor is easy to provide. For example, one's natural first thought (in a strategic communication setting) is that one's opponents will tell the truth, and a natural first thought (in the context of guessing games) is that all possible guesses are equally likely. In contrast, it is not at all clear how the level 0 ought to be specified in the setting of auctions. Crawford and Iriberri (2007) select what appear to be the two best proposals—that subjects either bid their valuation or randomize uniformly over the strategy space. However, it is not obvious whether either represents a natural first thought for how one's opponents will play. More to the point, it does not seem that there exist more plausible level-0 specifications than those proposed by Crawford and Iriberri (2007) (see sec. V for discussion). Thus, the failure of the level- k model in this context may well be down to the lack of a salient and psychologically plausible anchor.

Second, iterated reasoning is cognitively much more taxing in Bayesian games such as auctions than in the simpler settings to which the level- k model is normally applied. To calculate one's optimal action, one must combine one's conjecture about one's opponent's strategy with their distribution over types to arrive at their distribution over actions. One must then solve an (often nontrivial) optimization problem to find the optimal action given this distribution. Thus, even if a plausible level-0 anchor were available, it is unclear whether individuals would be able to correctly solve for the level-1 strategy—and even less clear whether they would be able to correctly identify the higher order strategies.

If this is right, it follows that the level- k model should be applied only in situations that possess both an intuitively appealing “first thought” along with easily computable best-response dynamics. Rigorously testing whether these features are indeed necessary for the predictive success of the model—along with identifying strategic settings that possess these apparently necessary features—would seem to be an important task for future research.

Data Availability

The replication package for this article (Rasooly 2022) can be found in the Harvard Dataverse: <https://doi.org/10.7910/DVN/SGHWJL>. The package includes all data used in the empirical sections (see the ReadMe file for details).

Appendix A

Proofs

A1. Proof of Proposition 1

Since our game has a symmetric and finite normal form, it must have a symmetric equilibrium (see Harsanyi [1967] and Cheng et al. [2004] for details). Turning to uniqueness, we now generalize three lemmas from Rasooly and Gavidia-Calderon (2020) to the case of mixed strategies. Let σ denote an arbitrary symmetric-equilibrium strategy, and for any $v \in \mathbb{X}$, let $s(v) \subseteq \mathbb{X}$ denote the set of bids that are submitted with positive probability at value v (i.e., the support at v).

LEMMA 1. In any symmetric equilibrium σ , $s(0) = \{0\}$.

Proof. Consider a player with $v = 0$. If they bid 0, they get (regardless of their opponents’ strategies) a payoff $\pi(v = 0, b = 0) = 0 \times \mathbb{P}(\text{win}|b = 0) - 0 = 0$. On the other hand, if they submit any bid $b \geq 1$, they get (given any opponent strategy profile) $\pi(v = 0, b) = 0 \times \mathbb{P}(\text{win}|b) - b = -b < 0$. Thus, bidding 0 strictly dominates bidding any $b \geq 1$, which means that $\mathbb{P}(b = 0|v = 0) = 1$ in any equilibrium (and in any symmetric equilibrium). QED

LEMMA 2 (monotonicity). Consider any symmetric equilibrium σ and consider any two values $v, v' \in \mathbb{X}$ with $v > v'$. Then for any $b \in s(v)$ and $b' \in s(v')$, $b \geq b'$.

Proof. For contradiction, let us suppose that $b < b'$. Since b is submitted with positive probability when the value is v , bidding b must be weakly optimal. In particular, it must be weakly better than bidding b' :

$$\begin{aligned} \pi(v, b) &\geq \pi(v, b') \\ \iff v\mathbb{P}(\text{win}|b) - b &\geq v\mathbb{P}(\text{win}|b') - b' && \text{(A1)} \\ \iff b' - b &\geq v(\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)). \end{aligned}$$

Similarly, since b' is submitted with positive probability when the value is v' , it must be weakly better than b :

$$\pi(v', b') \geq \pi(v', b)$$

$$\iff v' \mathbb{P}(\text{win}|b') - b' \geq v' \mathbb{P}(\text{win}|b) - b \quad (\text{A2})$$

$$\iff v'(\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)) \geq b' - b.$$

Inequalities (A1) and (A2) jointly imply that

$$v'(\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)) \geq v(\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b)). \quad (\text{A3})$$

Since $b' > b$, and b is a symmetric-equilibrium bid, it must be that $\mathbb{P}(\text{win}|b') > \mathbb{P}(\text{win}|b)$ (for example, one wins with a bid of b' , but not with a bid of b , if all of one's opponents bid b —and this happens with positive probability). Since $\mathbb{P}(\text{win}|b') - \mathbb{P}(\text{win}|b) > 0$, inequality (A3) then yields $v' \geq v$, which contradicts our initial assumption that $v' < v$. This establishes that $b \geq b'$ as claimed. QED

LEMMA 3 (no gaps). In any symmetric equilibrium σ , the bids that are submitted with positive probability are a set of consecutive integers.

Proof. Suppose for contradiction that this were false. Then there must exist two bids $b_{\text{high}} \geq b_{\text{low}} + 2$ that are submitted with positive probability even though no bid in between b_{high} and b_{low} is submitted with positive probability. Now if types were to deviate from bidding b_{high} to bidding $b_{\text{low}} + 1$, their payment would fall. However, their probability of winning would remain unchanged: both before and after the deviation, they win the auction if and only if all opponents bid b_{low} or lower. So this deviation would be strictly profitable, implying that this could not have been a symmetric equilibrium. QED

We now introduce the concept of jump form.

DEFINITION A1. For every bid $i \in B^s$, we define the i th jump j_i as

$$j_i = v_i + \mathbb{P}(b < i|v_i),$$

where v_i is the minimum value $v \in \mathbb{X}$ such that $\mathbb{P}(b = i|v) > 0$. We refer to the vector of jumps $j = (j_0, j_1, \dots, j_m)$ as a jump vector and say that it is increasing if $j_{i+1} > j_i$ for every $i \in \{0, 1, \dots, m-1\}$.

Under the lemmas outlined previously, this is an equivalent representation of behavioral strategies.

LEMMA 4. Any gapless and monotone strategy that satisfies $\mathbb{P}(b = 0|v = 0) = 1$ may be converted into an increasing jump vector j that satisfies $j_0 = 0$ and $j_i \geq 1$. Conversely, any increasing jump vector j that satisfies $j_0 = 0$ and $j_i \geq 1$ may be converted into a gapless and monotone strategy that satisfies $\mathbb{P}(b = 0|v = 0) = 1$.

Proof. To establish the first claim, take any gapless and monotone strategy that satisfies $\mathbb{P}(b = 0|v = 0) = 1$. This may be converted into a jump vector by computing each of the jump points: for every bid i in the support of the strategy, we identify the minimum value $v_i \in \mathbb{X}$ such that $\mathbb{P}(b = i|v_i) > 0$, and set $j_i = v_i + \mathbb{P}(b < i|v_i)$. Given that $\mathbb{P}(b = 0|v = 0) = 1$, and so $\mathbb{P}(b = 0|v = 0) > 0$, we know that $v_0 = 0$. Thus, $j_0 = v_0 + \mathbb{P}(b < 0|v_0) = 0$. Also, given $\mathbb{P}(b = 0|v = 0) = 1$, $v_1 \geq 1$, so $j_1 = v_1 + \mathbb{P}(b < 1|v_1) \geq 1$ as claimed.

To show that j is increasing, consider any two consecutive jump points $j_i, j_{i+1} \in j$. If the strategy is monotone, every value that submits a bid $b = i + 1$ must be weakly larger than every value that submits a bid $b = i$. So in particular, $v_{i+1} \geq v_i$.

This yields two possibilities: either $v_{i+1} > v_i$ or $v_{i+1} = v_i$. In the first instance ($v_{i+1} > v_i$),

$$\begin{aligned}
 j_{i+1} &\equiv v_{i+1} + \mathbb{P}(b < i + 1 | v_{i+1}) \\
 &\geq v_i + 1 + \mathbb{P}(b < i + 1 | v_{i+1}) \\
 &\geq v_i + 1 \\
 &> v_i + \mathbb{P}(b < i | v_i) \equiv j_i,
 \end{aligned}
 \tag{A4}$$

where the final inequality holds since $\mathbb{P}(b < i | v_i) < 1$ (which in turn holds since necessarily $\mathbb{P}(b = i | v_i) > 0$). So in that case, $j_{i+1} > j_i$. In the second instance ($v_{i+1} = v_i$), we have

$$\begin{aligned}
 j_{i+1} &\equiv v_{i+1} + \mathbb{P}(b < i + 1 | v_{i+1}) \\
 &= v_i + \mathbb{P}(b < i + 1 | v_i) \\
 &= v_i + \mathbb{P}(b < i | v_i) + \mathbb{P}(b = i | v_i) \\
 &> v_i + \mathbb{P}(b < i | v_i) \equiv j_i,
 \end{aligned}
 \tag{A5}$$

where the inequality again holds since $\mathbb{P}(b = i | v_i) > 0$. Either way, then, monotonicity implies that $j_{i+1} > j_i$. Since this true for all i , this then implies that that vector of jumps is increasing. So every monotone behavioral strategy generates an increasing vector of jumps.

To establish the second claim, take any increasing jump vector j that satisfies $j_0 = 0$ and $j_i \geq 1$. To begin, we show that these restrictions imply $\mathbb{P}(b = 0 | v = 0) = 1$. To see this, observe that, since $j_0 = 0$ and $j_0 = v_0 + \mathbb{P}(b < 0 | v_0) = v_0$, we have $v_0 = 0$. Also, $v_1 = \lfloor j_1 \rfloor \geq 1$ since $j_1 \geq 1$. Since $v_0 = 0$ and $v_1 \geq 1$, it follows that $\mathbb{P}(b = 0 | v = 0) = 1$ as claimed.

We now show how j uniquely pins down the rest of the behavioral strategy. Take j_1 . Since $v_1 = \lfloor j_1 \rfloor$ is the smallest value that bids 1, and the assumed strategy is monotone, any values $v < v_1$ must bid 0 with probability 1. Moreover, $\mathbb{P}(b = 0 | v_1) = j_1 - v_1$ (which may equal zero). Thus, j_1 uniquely determines which values bid 0. Next, take j_2 . Since $v_2 = \lfloor j_2 \rfloor$ is the smallest value that bids 2, and the assumed strategy is monotone, any values $v < v_2$ must bid less than 2. Given the previous, this means that any values $v \in \{v_1 + 1, \dots, v_2 - 1\}$ must bid 1 with probability 1. If $v_1 = v_2$, then $\mathbb{P}(b = 1 | v_2) = j_2 - j_1$; otherwise, $\mathbb{P}(b = 1 | v_2) = j_2 - v_2$. Thus, j_2 uniquely determines which values bid 1. Continuing inductively in this fashion, we see that j is consistent with a unique monotone and gapless bidding strategy. QED

Having recast strategies in jump form, we can now define our algorithm:

DEFINITION A2. Let $\hat{j} = (\hat{j}_1, \dots, \hat{j}_m)$ denote the output of the following algorithm:

1. Impose $\hat{j}_0 = 0$.
2. Starting at $i = 1$, find the minimum $j_i \in (\hat{j}_{i-1}, S]$ such that $\pi(v = \lfloor j_i \rfloor, b = i) - \pi(v = \lfloor j_i \rfloor, b = i - 1) \geq 0$.
3. If there is no such j_i , the algorithm terminates.

4. There does exist such a j_i , set $\hat{j}_i = j_i$ and repeat (from step 2) for $i + 1$.

We then observe the following.

LEMMA 5. The vector \hat{j} is the only possible symmetric equilibrium of the all-pay auction.

Proof. Suppose for contradiction that there were some symmetric equilibrium $j \neq \hat{j}$. If $j \neq \hat{j}$, there must be some jump at which they differ. Let j_i denote the first such jump (i.e., i is the smallest number such that $j \neq \hat{j}$). There are two possibilities: $j_i < \hat{j}_i$ or instead $j_i > \hat{j}_i$. We will argue that each leads to a contradiction.

To see why the first case is impossible, first define $\pi^i(v, b)$ and $\pi^j(v, b)$ as the payoff of a player with a value v who bids b given that their opponents all bid according to \hat{j} and j , respectively. Now recall that, by construction, \hat{j}_i is the minimum jump such that

$$\pi^i(v = \lfloor \hat{j}_i \rfloor, b = i) \geq \pi^i(v = \lfloor \hat{j}_i \rfloor, b = i - 1). \quad (\text{A6})$$

Since $j_i < \hat{j}_i$, this means that

$$\pi^i(v = \lfloor j_i \rfloor, b = i) < \pi^i(v = \lfloor j_i \rfloor, b = i - 1). \quad (\text{A7})$$

But then, given that one's opponents bid according to j , one should never bid $b = i$ with a value $v = \lfloor j_i \rfloor$. That is, $\mathbb{P}(b = i | v = \lfloor j_i \rfloor) = 0$. But then $v = \lfloor j_i \rfloor$ cannot be the smallest value that bids i (contradiction).

To see why the second case is impossible, recall again that every \hat{j}_i satisfies

$$\begin{aligned} \pi^i(v = \lfloor \hat{j}_i \rfloor, b = i) &\geq \pi^i(v = \lfloor \hat{j}_i \rfloor, b = i - 1) \\ \iff \lfloor \hat{j}_i \rfloor \mathbb{P}^i(\text{win} | b = i) - 1 &\geq \lfloor \hat{j}_i \rfloor \mathbb{P}^i(\text{win} | b = i - 1). \end{aligned} \quad (\text{A8})$$

As a preliminary, note that since the right-hand side is nonnegative, $\lfloor \hat{j}_i \rfloor \mathbb{P}^i(\text{win} | b = i) - 1 \geq 0$ and so $\lfloor \hat{j}_i \rfloor \geq 1$ for every $i \geq 1$. Let us record this fact.

Turning to the main argument, fix the value at $\lfloor \hat{j}_i \rfloor$ but consider now changing opponent bidding from \hat{j} to j . If $j_i > \hat{j}_i$, then $\mathbb{P}^j(\text{win} | b = i) > \mathbb{P}^{\hat{j}}(\text{win} | b = i)$. However, since j and \hat{j} agree for all previous jumps, $\mathbb{P}^j(\text{win} | b = i - 1) = \mathbb{P}^{\hat{j}}(\text{win} | b = i - 1)$. Since $\lfloor \hat{j}_i \rfloor \geq 1$ (established earlier), this all means that

$$\begin{aligned} \lfloor \hat{j}_i \rfloor \mathbb{P}^j(\text{win} | b = i) - 1 &> \lfloor \hat{j}_i \rfloor \mathbb{P}^j(\text{win} | b = i - 1) \\ \iff \pi^i(v = \lfloor \hat{j}_i \rfloor, b = i) &> \pi^i(v = \lfloor \hat{j}_i \rfloor, b = i - 1). \end{aligned} \quad (\text{A9})$$

If j is a symmetric equilibrium, this implies that $\mathbb{P}^j(b = i - 1 | v = \lfloor \hat{j}_i \rfloor) = 0$. However, this is impossible: since $j_i > \hat{j}_i$, and j and \hat{j} agree for previous jumps, j must allocate more probability mass to bids of $i - 1$, which means that $\mathbb{P}^j(b = i - 1 | v = \lfloor \hat{j}_i \rfloor) > \mathbb{P}^{\hat{j}}(b = i - 1 | v = \lfloor \hat{j}_i \rfloor) \geq 0$. QED

By the argument above, there is no possible symmetric equilibrium except for \hat{j} . As we noted earlier, however, our game must possess some symmetric equilibrium. From this, it follows that it has exactly one symmetric equilibrium (namely, \hat{j}). This concludes the proof. QED

A2. *Proof of Proposition 2*

See main text.

A3. *Proof of Proposition 3*

A proof of this result can be easily reconstructed from the proof of proposition 1 (and is available from the author upon request).

A4. *Proof of Proposition 4*

We argue by induction. First, we show that the desired pattern of bidding holds when $k = 1$. Next, we show that if it holds for any level $k - 1 \in \mathbb{K}$, then it also holds (provided that x is sufficiently large) for any level $k \in \mathbb{K}$.

Starting with the base case ($k = 1$), note that level-1 bidders choose a bid $b \in \mathbb{X}$ to maximize

$$\pi^1(v, b) = (v - b) \left(\frac{n - 1}{n} + \frac{b}{n(x + 1)} \right)^{n-1} \tag{A10}$$

(inserting $p = 1/n$ into eq. [6]). Considering this function's continuous extension, observe that

$$\frac{\partial \ln[\pi^1(v, b)]}{\partial b} = \frac{(n - 1)[v - (x + 1)] - nb}{(v - b)[b + (n - 1)(x + 1)]}. \tag{A11}$$

Since we are searching for the optimal bid, we can just consider the range $b \leq v$. In this range, the denominator of (A11) is nonnegative. Moreover, since $v \leq x$ for all $v \in \mathbb{X}$, the numerator is negative. Thus, $\partial \ln[\pi^1(v, b)]/\partial b \leq 0$ and so $\partial \pi^1(v, b)/\partial b \leq 0$. Since this is true of the function's continuous extension, it follows that $\pi(v, 0) \geq \pi(v, b)$ for every value $v \in \mathbb{X}$ and every positive integer bid $b \in \mathbb{X}$. That is, bidding zero is optimal (as claimed).

Next, we show that if the desired pattern of bidding holds for level $k - 1 \in \mathbb{K}$, then it also holds (provided that x is sufficiently large) for level $k \in \mathbb{K}$. Suppose then that all players of level $k - 1$ bid as proposed. Then they bid 0 when $v \leq v^*(k - 1)$, that is, at $\lfloor v^*(k - 1) + 1 \rfloor$ distinct values, and otherwise, they bid $k - 2$. Hence, for a level- k player,

$$\mathbb{P}^k(\text{win}|b) = \begin{cases} 1 & \text{if } b > k - 2, \\ \left[1 - p + p \left(\frac{\lfloor v^*(k - 1) + 1 \rfloor}{x + 1} \right) \right]^{n-1} & \text{if } b \in \{1, \dots, k - 2\}, \\ (1 - p)^{n-1} & \text{if } b = 0. \end{cases} \tag{A12}$$

Given this, it is immediate that $b^* \in \{0, 1, k - 1\}$; that is, there are at most three optimal bids. Moreover, $\pi^k(v, b = 0) - \pi^k(v, b = 1) \rightarrow v(1 - p)^{n-1} - (v - 1)(1 - p)^{n-1} < 0$ as $x \rightarrow \infty$, and so $b^* \in \{0, k - 1\}$ if x is sufficiently large. Finally, observe that the latter is optimal iff

$$\begin{aligned} \pi(v, b = k - 1) > \pi(v, b = 0) &\iff v - (k - 1) > v(1 - p)^{n-1} \\ &\iff v > \frac{k - 1}{1 - (1 - p)^{n-1}}, \end{aligned} \tag{A13}$$

that is, iff $v > v^*(k)$ (as claimed). QED

Appendix B

Illustrating the Algorithm

In this section, we illustrate the algorithm using two examples. In both, values are uniform so (in any symmetric strategy profile)

$$\begin{aligned} \mathbb{P}(b < i) &= \mathbb{P}(v < \lfloor j_i \rfloor) + \mathbb{P}(v = \lfloor j_i \rfloor)(j_i - \lfloor j_i \rfloor) \\ &= \frac{\lfloor j_i \rfloor}{S} + \frac{1}{S}(j_i - \lfloor j_i \rfloor) \\ &= \frac{j_i}{S}, \end{aligned} \tag{B1}$$

where S is the number of possible valuations.

EXAMPLE 1. Consider an all-pay auction with uniform values and $n = 2$. To find the first jump point, we look for the minimum $j_1 \in [2, S]$ such that

$$\pi(v = \lfloor j_1 \rfloor, b = 1) - \pi(v = \lfloor j_1 \rfloor, b = 0) \geq 0. \tag{B2}$$

Clearly, $\pi(v = \lfloor j_1 \rfloor, b = 0) = 0$ since $\mathbb{P}(\text{win} | b = 0) = 0$. In addition, $\mathbb{P}(\text{win} | b = 1) = \mathbb{P}(b = 0) = j_1/S$ and so $\pi(v = \lfloor j_1 \rfloor, b = 1) = \lfloor j_1 \rfloor j_1/S - 1$. So we look for the minimum j_1 such that

$$\frac{\lfloor j_1 \rfloor j_1}{S} - 1 \geq 0. \tag{B3}$$

Of course, the solution will depend on S . If $S = 100$, $j_1 = 10$ (and similarly $j_1 = \sqrt{S}$ for any square S). If $S = 99$, the solution remains at $j_1 = 10$. However, if $S = 101$, the solution becomes $j_1 = 10.1$, implying randomization. (One can prove that, as $S \rightarrow \infty$, the proportion of solutions that are integer solutions converges to $1/2$.) Let us suppose that $S = 101$ so that the first jump is $j_1 = 10.1$. To find the next jump, we look for the smallest j_2 that satisfies

$$\begin{aligned} \pi(v = \lfloor j_2 \rfloor, b = 2) - \pi(v = \lfloor j_2 \rfloor, b = 1) &\geq 0 \\ \iff \lfloor j_2 \rfloor [\mathbb{P}(\text{win} | b = 2) - \mathbb{P}(\text{win} | b = 1)] &\geq 1 \\ \iff \lfloor j_2 \rfloor [\mathbb{P}(b \leq 1) - \mathbb{P}(b = 0)] &\geq 1 \\ \iff \lfloor j_2 \rfloor \left(\frac{j_2}{101} - \frac{10.1}{101} \right) &\geq 1. \end{aligned} \tag{B4}$$

One can check that the solution is $j_2 = 16.4125$, so the equilibrium vector of jumps starts with $(0, 10.1, 16.4125, \dots)$. Returning to behavioral strategies, this means that all values $v \in \{0, 1, \dots, 9\}$ bid 0 (with probability 1), value $v = 10$ randomizes between $b = 0$ and $b = 1$ with probabilities .1 and .9, values $v \in \{11, \dots, 15\}$ bid 1 (with probability 1), and finally the value $v = 16$ bids $b = 1$ with probability .4125. (Since we have not computed j_3 , we cannot technically determine the probability with which value $v = 16$ bids $v = 2$ but in fact $j_3 \geq 17$ so this probability is $1 - .4125$.)

EXAMPLE 2. Consider now the first-price auction with canceled bids. To find the first jump j_i , we look for the smallest $j_i \in (0, S]$ such that

$$\begin{aligned} \pi(v = \lfloor j_i \rfloor, b = 1) &\geq \pi(v = \lfloor j_i \rfloor, b = 0) \\ \iff (\lfloor j_i \rfloor - 1)\mathbb{P}(\text{win} | b = 1) &\geq \lfloor j_i \rfloor \mathbb{P}(\text{win} | b = 0) \\ \iff (\lfloor j_i \rfloor - 1) \left(1 - p + p \frac{j_i}{S}\right)^{n-1} p &\geq \lfloor j_i \rfloor (1 - p)^{n-1} p \tag{B5} \\ \iff (\lfloor j_i \rfloor - 1) \left(1 - p + p \frac{j_i}{S}\right)^{n-1} &\geq \lfloor j_i \rfloor (1 - p)^{n-1}. \end{aligned}$$

For instance, suppose that $p = 1/2$, $n = 2$, and $S = 101$. Then (B5) reduces to

$$\lfloor j_i \rfloor j_i - 101 - j_i \geq 0, \tag{B6}$$

which is almost quadratic in j_i . One can check that this inequality is satisfied by $j_i = 11$; then $\lfloor j_i \rfloor j_i - 101 - j_i = 11^2 - 101 - 11 = 9 > 0$. On the other hand, it will not hold if $j_i < 11$, for then $\lfloor j_i \rfloor j_i - 101 - j_i \leq 10 \times 11 - 101 - 11 = -2 < 0$. So the minimum j_i that satisfies the inequality is 11; this is the first jump. This means that all values smaller than 11 must bid zero (with probability 1), but the value $v = 11$ never bids zero with any probability.

Appendix C

Optimizing p

To find the probability $p^* \in [0, 1]$ that maximizes $d(p)$, recall that $p^* > 0$ (from proposition 5). Moreover, it is obvious that $p^* < 1$: if $p = 1$, then $d(p) = 0$ (whereas $d(p) > 0$ is plainly possible). Hence, $p^* \in (0, 1)$, which means that the optimal p must satisfy the first-order condition

$$\frac{\partial}{\partial p} \int_0^x |\beta(v) - \beta^1(v)| dv = 0. \tag{C1}$$

One can check that $\beta(v) \geq \beta^1(v)$ for all $v \in [0, x]$. So (C1) is equivalent to

$$\frac{\partial}{\partial p} \int_0^x \beta(v) dv = \frac{\partial}{\partial p} \int_0^x \beta^1(v) dv. \tag{C2}$$

When $n = 2$, standard calculations reveal that

$$\begin{aligned}
\frac{\partial}{\partial p} \int_0^x \beta(v) dv &= \frac{\partial}{\partial p} \int_0^x \frac{pv^2}{2(pv - p + 1)} dv \\
&= \frac{\partial}{\partial p} \frac{x^2 [p(3p - 2) - 2(p - 1)^2 \ln(1 - p)]}{4p^2} \\
&= \frac{x^2 [(2 - p)p - 2(p - 1) \ln(1 - p)]}{2p^3}.
\end{aligned} \tag{C3}$$

Similarly, when $n \geq 3$, we find that

$$\begin{aligned}
\frac{\partial}{\partial p} \int_0^x \beta(v) dv &= \frac{\partial}{\partial p} \int_0^x \left(\frac{n-1}{n} \right) v - \frac{x(1-p)}{np} \left[1 - \left(\frac{1-p}{1-p+p(v/x)} \right)^{n-1} \right] dv \\
&= \frac{\partial}{\partial p} \frac{x^2 (2 + np[(n-1)p - 2]) - 2(1-p)^n}{2(n-2)np^2} \\
&= \frac{x^2 (p[n(1-p)^n - 2(1-p)^n + n + 2] + 2[(1-p)^n - 1] - np^2)}{(n-2)n(p-1)p^3}.
\end{aligned} \tag{C4}$$

Turning to the level-1 bidding function, we find that

$$\begin{aligned}
\frac{\partial}{\partial p} \int_0^x \beta^1(v) dv &= \frac{\partial}{\partial p} \int_{v^*}^x \left(\frac{n-1}{n} \right) v - \left(\frac{1-p}{p} \right) \frac{x}{n} dv \\
&= \frac{\partial}{\partial p} \frac{x^2 (np - 1)^2}{2(n-1)np^2} \\
&= \frac{x^2 (np - 1)}{(n-1)np^3},
\end{aligned} \tag{C5}$$

where $v^* = (1-p)x/(n-1)p$.

To find the optimal probability when $n = 2$, we now substitute (C3) and (C5) into the first-order condition (C2), obtaining

$$\frac{(2-p)p - 2(p-1) \ln(1-p)}{2p^3} = \frac{np - 1}{(n-1)np^3}. \tag{C6}$$

Similarly, when $n \geq 3$, we substitute (C4) and (C5) and into the first-order condition, yielding

$$\frac{p[n(1-p)^n - 2(1-p)^n + n + 2] + 2[(1-p)^n - 1] - np^2}{(n-2)n(p-1)p^3} = \frac{np - 1}{(n-1)np^3}. \tag{C7}$$

Either way, we obtain an equation that can be solved numerically to reveal candidates for p^* . It is then a straightforward matter to check which of these candidates generates a larger distance. Since our problem must have a global maximizer (by the Weierstrass theorem), and the global maximizer must satisfy the first-order condition, the p satisfying the first-order condition that generates the largest distance must be the global maximizer.

TABLE C1
OPTIMAL CANCELLATION PROBABILITIES (Rounded)

	<i>n</i>						
	2	3	4	5	6	7	8
<i>p</i> *	.536	.343	.256	.204	.170	.145	.127

Table C1 outlines the solutions for $n \in \{1, \dots, 8\}$. As can be seen, the solution is generally quite close to $1/n$ —so the bound identified in proposition 5 comes close to identifying the exact solution. Note also that, as might be expected, the optimal probability p does not depend on the scale parameter x .

Appendix D

Tables and Figures

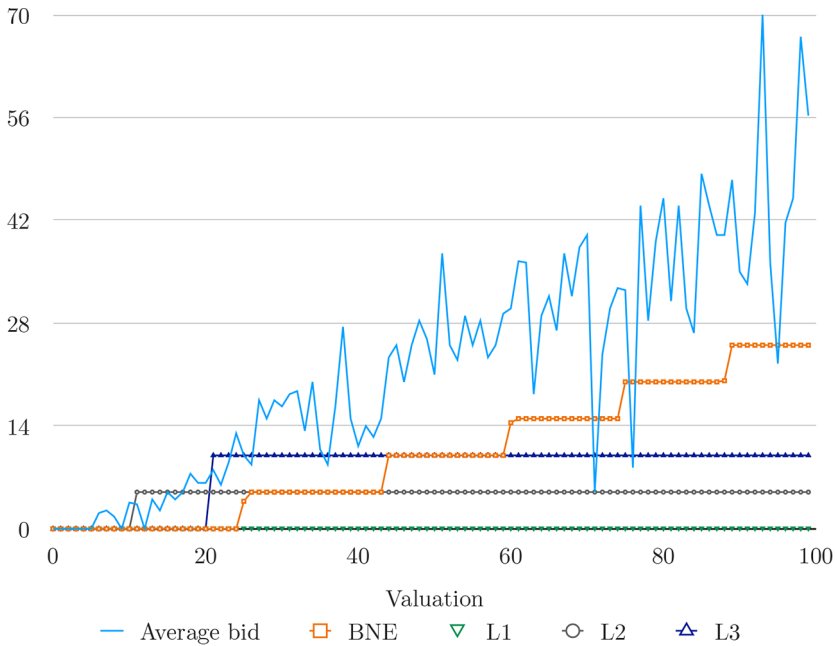


FIG. D1.—The first-price auction when bids are multiples of 5.




FIG. D2.—The all-pay auction when bids are multiples of 5.

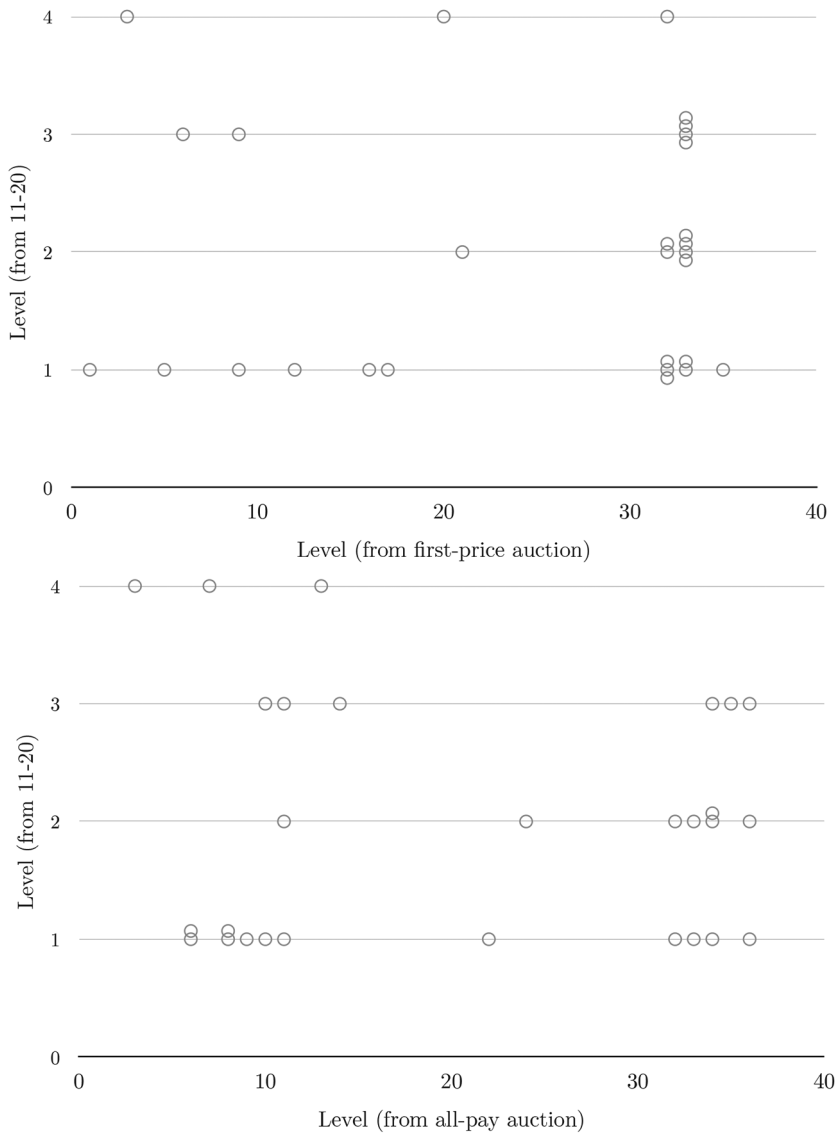


FIG. D3.—Comparing estimated levels. The top panel plots levels as estimated from the first-price auction against levels as inferred from the 11–20 game. The estimated correlation between these variables is 0 (95% confidence interval: -0.37 to 0.37). The bottom panel plots levels as estimated from the all-pay auction against levels as inferred from the 11–20 game. The estimated correlation between these variables is -0.09 (95% confidence interval: -0.44 to 0.30).

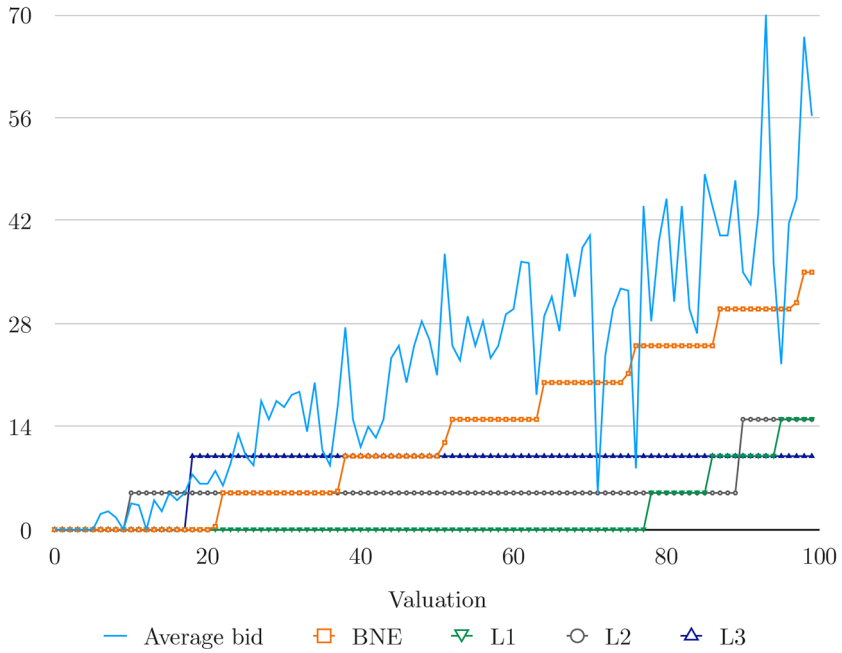


FIG. D4.—Risk aversion (multiples of 5 treatment).

TABLE D1
BALANCE TABLE

	Integer Bids	Multiples of 5	<i>p</i> -Value
Mean age	31.4	32.0	.29
Share male	.54	.43	.33
Share mathematics	.39	.38	.95
Share humanities	.24	.14	.36
Share social sciences	.30	.29	.93

NOTE.—This table displays how demographics (age, sex, and subject) vary between the two treatment groups. The column “*p*-value” reports *t*-tests of the hypothesis that the relevant variable is on average equal across the groups.

TABLE D2
AVERAGE BIDS

	FIRST PRICE			ALL PAY		
	T1	T2	<i>p</i> -Value	T1	T2	<i>p</i> -Value
Average bid	23.6	23.0	.85	22.2	25.4	.12
Average value	50.9	49.6	.27	49.4	50.8	.89
Ratio	.46	.46		.45	.50	

NOTE.—This table reports the average bids across two treatments and auction structures. The column “*p*-value” reports *t*-tests of the hypothesis that the relevant variable is on average equal across the groups.

TABLE D3
COMPARING EQUILIBRIUM AND LEVEL *k*

		T1 FP	T1 AP	T2 FP	T2 AP
BNE	LL	-4,407.9	-4,332.7	-1,024.7	-1,042.2
	BIC	8,820.5	8,670.2	2,053.1	2,088.1
L1	LL	-4,530.4	-4,501.2	-1,088.8	-1,128.6
	BIC	9,065.4	9,007.1	2,181.3	2,260.9
L1-L2	LL	-4,530.4	-4,501.2	-1,085	-1,126.9
	BIC	9,070.1	9,011.7	2,177.4	2,261.2
L1-L3	LL	-4,530.1	-4,501.2	-1,071.1	-1,120.7
	BIC	9,074.3	9,016.4	2,153.4	2,252.6

NOTE.—This table reports the log likelihoods and associated Bayes information criteria values for four different structural models. In the first model, equilibrium is the only type. The subsequent models are populated by level 1 types, level 1 and 2 types, and level 1, 2, and 3 types, respectively.

TABLE D4
THE HYBRID MODEL

	T1 FP		T1 AP		T2 FP		T2 AP	
	Subject	Round	Subject	Round	Subject	Round	Subject	Round
ρ_0	.758 (.059)	.760 (.046)	.815 (.053)	.757 (.06)	.905 (.066)	.858 (.054)	.741 (.149)	.788 (.083)
ρ_1	.036 (.025)	.046 (.02)	.000 (.0)	.010 (.01)	.095 (.066)	.044 (.034)	.048 (.048)	.071 (.04)
ρ_2	.133 (.051)	.027 (.016)	.019 (.019)	.055 (.022)	.000 (.0)	.000 (.0)	.048 (.048)	.024 (.024)
ρ_3	.073 (.041)	.167 (.041)	.167 (.051)	.178 (.057)	.000 (.0)	.098 (.048)	.164 (.141)	.117 (.073)
σ_0	27.3 (1.9)	27.3 (1.8)	19.9 (1.6)	20.2 (1.5)	23.2 (3.0)	23.7 (2.4)	21.7 (2.8)	21.3 (2.0)
σ_1	3.8 (2.0)	1.8 (.3)		2.3 (.4)	5.4 (2.2)	6.3 (.7)	2.7 (.0)	2.4 (.4)
σ_2	9.2 (.6)	4.2 (.5)	1.1 (.0)	1.0 (.0)			2.7 (.0)	1.0 (.0)
σ_3	13.3 (1.1)	10.5 (1.2)	5.3 (.6)	7.9 (2.1)		4.1 (.9)	14.5 (7.0)	6.5 (4.1)
LL	-4,241.5	-4,217.1	-4,043.8	-4,004.1	-987.9	-982.6	-962.9	-940.9

NOTE.—The “subject” columns report analyses in which subjects’ levels are fixed; the “round” columns report analyses in which levels are allowed to vary between rounds. The symbols ρ_i and σ_i respectively denote the probability and noise parameter of type $i \in \mathbb{K}$ (equilibrium is indexed by $i = 0$). LL abbreviates “log likelihood,” and jackknife standard errors are in parentheses. When $\rho_i = 0$, the associated parameter σ_i is unidentified and therefore omitted from the table.

TABLE D5
BIDS BY ESTIMATED LEVEL

	FIRST PRICE			ALL PAY		
	Average Bid	Average Value	Ratio	Average Bid	Average Value	Ratio
L1	19.8	48.8	.406	22.1	51.8	.427
L2	20.8	51.6	.403	21.4	47.8	.447
L3	25.7	51.2	.502	26.6	49.7	.536

NOTE.—This table reports the average bids and valuations of the subjects who respectively chose 19 (level 1), 18 (level 2), and 17 (level 3) in the 11–20 game.

TABLE D6
ROBUSTNESS CHECKS

		T1 FP	T1 AP	T2 FP	T2 AP
Risk aversion	BNE	19.4		17.9	
	Level <i>k</i>	26.9		22.7	
Cognitive hierarchy	BNE	21.7	18.2	19.7	18.8
	Level <i>k</i>	31.3	23.8	26.6	29.6
Dominated bids	BNE	22.6	19.4	19.5	19.0
	Level <i>k</i>	30.8	29.9	22.7	26.9
Tie breaking	BNE	21.7	18.2	19.7	18.8
	Level <i>k</i>	30.3	29.5	23.4	26.6
Dropping round 2	BNE	18.4	18.2	18.3	17.4
	Level <i>k</i>	26.5	29.0	22.2	24.0

NOTE.—This table reports the root-mean-square prediction errors of the relevant theories following a series of robustness checks. In every case, level-*k* predictions are obtained by assigning each data point the level from the 1–3 range that minimizes the model’s prediction error (a procedure that will tend to overstate the model’s predictive performance).

References

- Alaoui, L. and A. Penta. 2016. “Endogenous Depth of Reasoning.” *Rev. Econ. Studies* 83 (4): 1297–333.
- An, Y. 2017. “Identification of First-Price Auctions with Non-Equilibrium Beliefs: A Measurement Error Approach.” *J. Econometrics* 200 (2): 326–43.
- Arad, A., and A. Rubinstein. 2012. “The 11–20 Money Request Game: A Level-*k* Reasoning Study.” *A.E.R.* 102 (7): 3561–73.
- Arrow, K. J. 1970. *Essays in the Theory of Risk-Bearing*. Amsterdam: North-Holland.
- Bleichrodt, H., J. N. Doctor, Y. Gao, C. Li, D. Meeker, and P. P. Wakker. 2019. “Resolving Rabin’s Paradox.” *J. Risk and Uncertainty* 59 (3): 239–60.
- Bosch-Domènech, A., J. G. Montalvo, R. Nagel, and A. Satorra. 2002. “One, Two, (Three), Infinity, . . . : Newspaper and Lab Beauty-Contest Experiments.” *A.E.R.* 92 (5): 1687–701.
- Brocas, I., J. D. Carrillo, S. W. Wang, and C. F. Camerer. 2014. “Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games.” *Rev. Econ. Studies* 81 (3): 944–70.
- Cai, H., and J. T.-Y. Wang. 2006. “Overcommunication in Strategic Information Transmission Games.” *Games and Econ. Behavior* 56 (1): 7–36.
- Camerer, C. F., T.-H. Ho, and J.-K. Chong. 2004. “A Cognitive Hierarchy Model of Games.” *Q.J.E.* 119 (3): 861–98.

- Chen, D. L., M. Schonger, and C. Wickens. 2016. "oTree—an Open-Source Platform for Laboratory, Online, and Field Experiments." *J. Behavioral and Experimental Finance* 9:88–97.
- Cheng, S.-F., D. M. Reeves, Y. Vorobeychik, and M. P. Wellman. 2004. "Notes on Equilibria in Symmetric Games." In *Proceedings of the 6th International Workshop on Game Theoretic and Decision Theoretic Agents*, 71–78. New York: GTDT.
- Costa-Gomes, M. A., V. P. Crawford, and N. Iriberri. 2009. "Comparing Models of Strategic Thinking in Van Huyck, Battalio, and Beil's Coordination Games." *J. European Econ. Assoc.* 7 (2–3): 365–76.
- Costa-Gomes, M. A., and M. Shimoji. 2015. "A Comment on 'Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study.'" *Econometrica* 83 (1): 375–83.
- Crawford, V. P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *A.E.R.* 93 (1): 133–49.
- Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *J. Econ. Literature* 51 (1): 5–62.
- Crawford, V. P., and N. Iriberri. 2007. "Level- k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica* 75 (6): 1721–70.
- Crawford, V. P., T. Kugler, Z. Neeman, and A. Pauzner. 2009. "Behaviorally Optimal Auction Design: Examples and Observations." *J. European Econ. Assoc.* 7 (2–3): 377–87.
- Crosetto, P., and A. Filippin. 2013. "The 'Bomb' Risk Elicitation Task." *J. Risk and Uncertainty* 47 (1): 31–65.
- De Clippel, G., R. Saran, and R. Serrano. 2019. "Level- k Mechanism Design." *Rev. Econ. Studies* 86 (3): 1207–27.
- Filiz-Ozbay, E., and E. Y. Ozbay. 2007. "Auctions with Anticipated Regret: Theory and Experiment." *A.E.R.* 97 (4): 1407–18.
- Galavotti, S., L. Moretti, and P. Valbonesi. 2018. "Sophisticated Bidders in Beauty-Contest Auctions." *American Econ. J. Microeconomics* 10 (4): 1–26.
- Georganas, S., P. J. Healy, and R. A. Weber. 2015. "On the Persistence of Strategic Sophistication." *J. Econ. Theory* 159: 369–400.
- Gillen, B. 2009. "Identification and Estimation of Level- k Auctions." <https://ssrn.com/abstract=1337843>.
- Harsanyi, J. C. 1967. "Games with Incomplete Information Played by 'Bayesian' Players, I–III Part I. The Basic Model." *Management Sci.* 14 (3): 159–82.
- Harstad, R. M., J. H. Kagel, and D. Levin. 1990. "Equilibrium Bid Functions for Auctions with an Uncertain Number of Bidders." *Econ. Letters* 33 (1): 35–40.
- Holt, C. A. 1980. "Competitive Bidding for Contracts under Alternative Auction Procedures." *J.P.E.* 88 (3): 433–45.
- Hortaçsu, A., F. Luco, S. L. Puller, and D. Zhu. 2019. "Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets." *A.E.R.* 109 (12): 4302–42.
- Ivanov, A., D. Levin, and M. Niederle. 2010. "Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study." *Econometrica* 78 (4): 1435–52.
- Kagel, J., and D. Levin. 1995. "Auctions: A Survey of Experimental Research." In *Handbook of Experimental Economics*, edited by J. Kagel and A. Roth. Princeton, NJ: Princeton Univ. Press.
- Kirchkamp, O., and J. P. Reiß. 2011. "Out-of-Equilibrium Bids in First-Price Auctions: Wrong Expectations or Wrong Bids." *Econ. J.* 121 (557): 1361–97.

- Klemperer, P. 1999. "Auction Theory: A Guide to the Literature." *J. Econ. Surveys* 13(3): 227–86.
- Krishna, V. 2009. *Auction Theory*. Burlington, MA: Academic Press.
- Nagel, R. 1995. "Unraveling in Guessing Games: An Experimental Study." *A.E.R.* 85 (5): 1313–26.
- Noussair, C., and J. Silver. 2006. "Behavior in All-Pay Auctions with Incomplete Information." *Games and Econ. Behavior* 55 (1): 189–206.
- Rabin, M. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica* 68 (5): 1281–92.
- Rasooly, I. 2022. "Replication Data for 'Going. . . Going. . . Wrong: A Test of the Level- k (and Cognitive-Hierarchy) Models of Bidding Behavior.'" Harvard Dataverse. <https://doi.org/10.7910/DVN/SGHWJL>.
- Rasooly, I., and C. Gavidia-Calderon. 2020. "The Importance of Being Discrete: On the Inaccuracy of Continuous Approximations in Auction Theory." <https://arxiv.org/abs/2006.03016>.
- Stahl, D. O., and P. W. Wilson. 1994. "Experimental Evidence on Players' Models of Other Players." *J. Econ. Behavior and Org.* 25 (3): 309–27.