

Level- k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?

Vincent P. Crawford and Nagore Iriberrí¹

University of California, San Diego

15 October 2005; revised 4 November 2005

"The common curse of mankind, folly and ignorance, be thine in great revenue!"

—William Shakespeare, *Troilus and Cressida*

Abstract: This paper proposes a structural non-equilibrium model of initial responses to incomplete-information games based on "level- k " thinking, which describes behavior in many experiments with complete-information games. We derive the model's implications in first- and second-price auctions with general information structures, compare them to equilibrium and Eyster and Rabin's (2005) "cursed equilibrium," and evaluate the model's potential to explain behavior in auction experiments. The level- k model generalizes many insights from equilibrium auction theory. It also allows a unified explanation of the winner's curse in common-value auctions and overbidding in those independent-private-value auctions without the uniform value distributions used in most experiments.

Keywords: common-value auctions, winner's curse, overbidding, bounded rationality, level- k model, non-equilibrium strategic thinking, behavioral game theory, experiments

JEL classification numbers: C72, C92, D44, D82

¹Email: vcrawfor@dss.ucsd.edu and niriberr@econ.ucsd.edu. We are grateful to the National Science Foundation (Crawford) and the Centro de Formacion del Banco de España (Iriberrí) for research support; to Pierpaolo Battigalli, Drew Fudenberg, Mark Isaac, Thomas Palfrey, Charles Plott, Jose Gonzalo Rangel, Ricardo Serrano-Padial, Joel Sobel, and especially Dan Levin for helpful discussions; and to John Kagel for sharing the data from the experiments of Kagel and Levin (1986, 2002), Garvin and Kagel (1994), and Avery and Kagel (1997). An Appendix with supporting calculations is posted, along with the paper, at <http://dss.ucsd.edu/~vcrawfor/#Auctions>.

1. Introduction

Common-value auctions, in which the value of the object being sold is unknown but the same to all bidders *ex post* and each bidder receives a private signal that is correlated with the value, have been studied intensively, both theoretically and empirically, since Milgrom and Weber (1982; "MW"); see the surveys by McAfee and McMillan (1987, Section X), Milgrom (1985, Section 4; 1987, Section 4), Wilson (1992, Section 4.2), and Klemperer (2000, Chapter 1).

A central problem in this area is explaining the "winner's curse," the frequent tendency for winning bidders in common-value auctions to find, on average, that they have bid more than the object being sold is worth.² The curse, as we shall call it, was first noted in oil-lease auctions by petroleum engineers (Capen, Clapp, and Campbell (1971)) and studied theoretically by Wilson (1969). It has since been detected in many analyses of field data (Hendricks, Porter, and Boudreau (1987); Hendricks and Porter (1988); and the papers surveyed in McAfee and McMillan (1987, Section XII), Thaler (1988), Wilson (1992, Section 9.2), and Laffont (1997, Section 3)). The curse has also been observed in laboratory experiments with precise control of the information conditions on which it depends (Bazerman and Samuelson (1983); Kagel and Levin (1986; "KL"); Dyer, Kagel, and Levin (1989); Lind and Plott (1991; "LP"); and the papers surveyed in Kagel (1995, Section II) and KL (2002)). Finally, curse-like phenomena have been observed in non-auction settings that share the informational features of common-value auctions: bilateral negotiations in the Acquiring a Company game in Samuelson and Bazerman (1985), Holt and Sherman (1994; "HS"), Tor and Bazerman (2003), and Charness and Levin (2005); the Monty Hall game in Friedman (1998), Tor and Bazerman (2003), and Palacios-Huerta (2003); zero-sum betting with asymmetric information in Sovik (2000) and Sonsino, Erev, and Gilat (2002); and voting and jury decisions in Feddersen and Pesendorfer (1996, 1997, 1998). There is also an experimental literature on independent-private-value auctions, which documents a widespread tendency for subjects to bid higher than in the risk-neutral Bayesian equilibrium—though not usually to the point of making losses as in common-value auctions; see Cox, Smith, and Walker (1983, 1988); Goeree, Holt, and Palfrey (2002; "GHP"); and the references cited there.

The curse is often attributed informally to bidders' failure to adjust their value estimates for the information revealed by winning. Such adjustments are illustrated by the symmetric Bayesian equilibrium of a first- or second-price auction with symmetric bidders, where bidders adjust their

²This definition is the one most often used in the literature. We use the term "common-value" to include MW's (1982) general notion of interdependent values with affiliated signals as well as pure common values.

expected values for the fact that the winner's private signal must have been more favorable than all others' signals, and so overestimates the value based on all available information.³ But despite the empirical importance of overbidding in independent-private-value auctions and curse-like phenomena in common-value auctions, there have been few attempts to model them formally.

KL (1986) and HS (1994) formalize the intuition behind the curse in models in which "naïve" bidders do not adjust their value estimates for the information revealed by winning, but otherwise follow equilibrium logic. Eyster and Rabin's (2002, 2005; "ER") notion of "cursed equilibrium" generalizes KL's and HS's models to allow intermediate levels of value adjustment, ranging from standard equilibrium with full adjustment to "fully-cursed" equilibrium with no adjustment. ER also generalize KL's and HS's models from auctions and bilateral exchange to other kinds of incomplete-information games.⁴ All three models allow players to deviate from equilibrium only to the extent that they do not draw correct inferences from the outcome. Thus their predictions coincide with equilibrium in games in which such inferences are not relevant, and they do not help to explain non-equilibrium behavior in independent-private-value auctions.

Other analyses, also assuming equilibrium, seek to explain overbidding in independent-private-value auctions via various deviations from risk-neutral expected-monetary-payoff maximization: risk aversion in Cox, Smith, and Walker (1983, 1988) and HS (2000); the "joy of winning" in Cox, Smith, and Walker (1992) and HS (1994); and both of these plus nonlinear probability weighting, using McKelvey and Palfrey's (1995) notion of quantal response equilibrium ("QRE"), in GHP (2002).⁵

These explanations all assume the perfect coordination of beliefs about others' strategies that is characteristic of equilibrium analysis. Such coordination is plausible when bidders have had

³A bidder's bid should be chosen as if it were certain to win because it affects the bidder's payoff only when it wins.

⁴In Samuelson and Bazerman's (1985) experiments with the Acquiring a Company game, both less- and more-informed subjects tend to choose as if their (more- or less-informed) partner's information was the same as their own. In Acquiring a Company, cursed equilibrium would assume this for a less-informed player but not a more-informed player. It is unclear how to extend Samuelson and Bazerman's interpretation to auctions in which each player has some private information, so that no one is unambiguously less- or more-informed. Neither of the obvious choices—that a player ignores his own private information, or alternatively assumes that all others share it—seems sensible.

⁵QRE is a generalization of equilibrium that allows players' choices to be noisy, with the probability of each choice increasing in its expected payoff, given the distribution of others' decisions; a QRE is thus a fixed point in the space of players' choice distributions. GHP (2002) use QRE to analyze the results of experiments with independent-private-value auctions. To our knowledge QRE has not been used to analyze common-value auctions. Risk aversion has been applied mainly to explain overbidding in independent-private-value auctions, with the notable exception of HS (2000). As LP (1991) note, common-value auctions with risk aversion are not well understood theoretically.

ample opportunity to learn from experience with analogous auctions.⁶ But some auctions that have been studied using field data lack enough clear precedents to make equilibrium a plausible hypothesis for initial responses; and subjects may learn slowly in auction experiments, especially with common values (LP (1991); Ball, Bazerman, and Carroll (1991); Garvin and Kagel (1994); Kagel and Richard (2001); and Palacios-Huerta (2003)). The justification for equilibrium then depends on strategic thinking rather than learning, but such thinking may not follow the fixed-point logic of equilibrium. It may then be just as plausible to relax the assumption of equilibrium as to relax correct value adjustment or risk-neutral expected-money-payoff maximization.

Progress via relaxing equilibrium depends on specifying a structural model that accurately describes initial responses to games.⁷ In this paper we reconsider the winner's curse in common-value auctions and overbidding in independent-private-value auctions using a non-equilibrium model of initial responses based on "level- k " thinking, introduced by Stahl and Wilson (1994, 1995) and Nagel (1995) and developed and further applied by Ho, Camerer, and Weigelt (1998); Costa-Gomes, Crawford, and Broseta (2001); Crawford (2003); Camerer, Ho, and Chong (2004; "CHC"); Costa-Gomes and Crawford (2004); and Crawford and Iriberry (2005). The level- k model has strong experimental support, which should help to allay the common concern that once one departs from equilibrium, "anything is possible." We focus on symmetric first- and second-price auctions, leaving their progressive Dutch and English counterparts for future work.

The level- k model has the potential to give a unified explanation of overbidding in independent-private-value and common-value auctions as well as curse-like phenomena in other settings. It also promises to establish a link between empirical auction studies and non-auction experiments on strategic thinking, and thereby to bring a large body of auction evidence to bear

⁶Such experience might justify fully cursed equilibrium, for instance, by teaching bidders the tradeoff between the cost of higher bids and their increased probability of winning without also teaching them to avoid the curse. In the field bidders seldom observe others' value estimates, which impedes learning about the curse. In most of the relevant experiments, subjects' bids and signals were made public after each round, but even experienced subjects may focus on the relationship between the winner's signal and bid and the realized value of the object, without looking for relationships like the curse. It seems much harder to justify less than fully-cursed equilibrium, because once one realizes there may be a relationship to look for, there is no obvious reason to stop at intermediate levels of cursedness.

⁷Maintaining common knowledge of rationality but otherwise leaving beliefs unrestricted yields notions like rationalizability, which implies some restrictions on behavior in first-price auctions or common-value second-price auctions, and duplicates equilibrium in independent-private-value second-price auctions. k -level rationalizability—consistency with rationality and mutual certainty of $(k-1)$ -level rationalizability—implies bounds on behavior in first-price auctions characterized in Battigalli and Siniscalchi (2003) and restricts behavior in common-value second-price auctions, and also duplicates equilibrium in independent-private-value second-price auctions. By contrast, our approach dispenses with common knowledge of rationality (and beliefs) but normally yields unique predictions.

on the issue of how best to model initial responses to games. Finally, it allows us to explore the robustness of equilibrium auction theory's conclusions to failures of the equilibrium assumption.

The level- k model allows behavior to be heterogeneous, but it assumes that each player's behavior is drawn from a common distribution over a particular hierarchy of decision rules or *types*. Type L_k for $k > 0$ anchors its beliefs in a nonstrategic L_0 type and adjusts them via thought-experiments with iterated best responses: L_1 best responds to L_0 , L_2 to L_1 , etc. L_1 and L_2 have accurate models of the game and are rational; they depart from equilibrium only in basing their beliefs on simplified models of other players.⁸ This yields a workable model of others' decisions while avoiding much of the cognitive complexity of equilibrium analysis.⁹ In applications the population type distribution is usually translated from previous work or estimated from the current dataset. The estimated distribution tends to be stable across games, with most of the weight on L_1 and L_2 . Thus the anchoring L_0 type exists mainly in the minds of higher types.

Even so, the specification of L_0 is the key to the model's explanatory power and the main issue that arises in extending the level- k model from complete- to incomplete-information games. We compare two specifications, both nonstrategic as is usual in level- k analyses. A *random* L_0 bids uniformly randomly over the feasible range, as in the complete-information level- k analyses of Stahl and Wilson (1994, 1995); Costa-Gomes, Crawford, and Broseta (2001); CHC (2004); and Costa-Gomes and Crawford (2004).¹⁰ A *truthful* L_0 bids the value that its own private information

⁸Charness and Levin (2005) conduct an interesting experimental test of "simplified models of others" explanations of the curse like the one proposed here, in an Acquiring a Company design with a "robot" treatment in which a single decision-maker faces an updating problem that is cognitively the same as the one that underlies the curse. They find that the curse persists in their treatment, and suggest that their results favor explanations based on limited cognition in Bayesian updating or understanding the problem rather than simplified models of others. Their results suggest that the curse is due to some form of limited cognition rather than strategic uncertainty; but their analysis leaves open the possibility that something like the level- k model accurately describes initial responses to environments, interactive or not, that pose cognitive difficulties that are isomorphic to those of predicting other players' strategic decisions.

⁹In Selten's (1998) words: "Basic concepts in game theory are often circular in the sense that they are based on definitions by implicit properties.... Boundedly...rational strategic reasoning seems to avoid circular concepts. It directly results in a procedure by which a problem solution is found. Each step of the procedure is simple, even if many case distinctions by simple criteria may have to be made." Costa-Gomes and Crawford (2004) summarize the evidence for the level- k model and give support for our assumptions that L_2 best responds to an L_1 without decision errors, unlike in Stahl and Wilson (1994, 1995); and to L_1 alone rather than a mixture of L_1 and L_0 , unlike *Worldly* in Stahl and Wilson (1995) and L_2 in CHC (2004). We confine attention to L_0 , L_1 , and L_2 because they well illustrate the model's potential to explain auction behavior and the evidence suggests that higher types are comparatively rare.

¹⁰In the only other incomplete-information level- k model of which we are aware, CHC (2004, Section VI.A) use their closely related "cognitive hierarchy" model, with a random L_0 , to explain curse-like phenomena in Sonsino, Erev, and Gilat's (2002) and Sovik's (2000) experimental results on zero-sum betting with asymmetric information. One could easily imagine other possible specifications of random L_0 , for example restricting the range of variation to bids that can be rational, given one's own information. But bearing in mind that L_0 is only the starting point for a player's analysis of others' behavior, simplicity is a virtue. Ultimately the "correct" specification is an empirical question.

suggests, taken by itself. Although truthfulness has no natural meaning in most settings for which level- k analyses have been conducted, in auctions it is both meaningful and plausible, given that $L0$ is only the starting point for a strategic analysis.¹¹ We call the $L1$ and $L2$ types based on a random $L0$, *random $L1$* and *$L2$* types, with analogous terms for the *truthful $L1$* and *$L2$* types based on a truthful $L0$. We stress that these $L1$ and $L2$ types need not be random or truthful themselves.

Although a level- k model's predictions coincide with equilibrium in many simple games, in games as complex as auctions they may deviate systematically from equilibrium. The deviations are determined by the same two factors that determine an equilibrium bidder's bidding strategy—*value adjustment* for the information revealed by winning and the *bidding trade-off* between a higher bid's cost and its increased probability of winning—but the influences of these factors are altered by a level- k type's non-equilibrium beliefs. The pattern of $L1$'s and $L2$'s deviations across first- and second-price common- and independent-private-value auctions determines the extent to which a level- k model with an empirically plausible type distribution allows a unified explanation of the systematic, heterogeneous deviations from equilibrium such auctions often evoke.

Our analysis yields two main conclusions. First, many of the insights of equilibrium auction theory extend, suitably interpreted, to level- k auction theory.¹² Second, a random level- k model (but not a truthful level- k model) can yield an empirically plausible explanation of both the winner's curse in common-value auctions and overbidding in those independent-private-value auctions without the uniform value distributions used in most experiments.

In common value auctions, because a random $L0$'s bids are independent of its private signal, a random $L1$ completely ignores the information revealed by winning, just as ER's fully-cursed equilibrium bidders do. In a second-price auction the bidding trade-off is neutral and a random $L1$'s failure to adjust the value makes its bids coincide with a fully-cursed equilibrium bidder's, so that it normally overbids relative to equilibrium. In a first-price auction a random $L1$ differs from

¹¹Our truthful $L0$ is equivalent to LP's (1991) "naive model" and our random $L1$ is close to their "private-value" model. Our truthful $L0$ is also reminiscent of the truthful sender type $W0$ in Crawford's (2003) level- k analysis of strategic deception via cheap talk, which also appears frequently in the informal literature on deception and receives some support in communication experiments (Crawford (1998), Cai and Wang (2005), and the references cited there). Models that adapt $L0$ to the setting in other ways include Ho, Camerer, and Weigelt's (1998) analysis of guessing games, where $L0$ is random with an estimated central tendency; and Crawford and Iriberti's (2005) analysis of Hide-and-Seek games, where $L0$ responds to the non-neutral framing of locations. By contrast, the level- k model's other main assumption, the adjustment of higher-level types' beliefs via iterated best responses, appears to allow a satisfactory account of initial responses to many different kinds of games.

¹²To the extent that equilibrium insights do not generalize, it is mainly because level- k types, by best responding to level- $(k-1)$ types, break the symmetry of a standard equilibrium analysis, which creates difficulties like those in equilibrium analyses of asymmetric auctions (McAfee and McMillan (1987, Section VII), Maskin and Riley (2000)).

a fully-cursed-equilibrium bidder in using its non-equilibrium beliefs to evaluate a non-neutral bidding trade-off. In general a random $L1$ overbids relative to equilibrium but may either overbid or underbid relative to or coincide with fully-cursed equilibrium. In the leading example that has been studied experimentally, KL's, a random $L1$ coincides with fully-cursed equilibrium with common values; and coincides with it and equilibrium with uniform, independent private values. Without uniformity, in general, a random $L1$ may underbid, overbid, or coincide with equilibrium. A random $L1$ coincides with equilibrium except for one possible valuation in one of GHP's (2002) independent-private-value treatments with discrete, slightly non-uniform values.

In a second- or first-price auction, a random $L1$'s bidding strategy is increasing in its signal, and so in common-value auctions a random $L2$ does adjust its value estimate for the information revealed by winning. Its value adjustment reflects the same logic as the second-price equilibrium bidding strategy, in that it bids according to the expected value given its own signal, conditional on *just* winning; but its non-equilibrium beliefs do not anticipate winning if and only if it has the highest signal, which leads to a different adjustment. In general, value adjustment tends to make bidders' bids strategic substitutes, because higher others' bids make a bidder believe that winning means others' signals are (stochastically) lower, which makes the curse seem worse and lowers the expected value conditional on winning, and so tends to lower the bidder's optimal bid. In a second-price auction, to the extent that a random $L1$ overbids relative to equilibrium, the strategic substitutability of value adjustment makes a random $L2$ underbid. This effect is also present in a first-price auction, but there the bidding trade-off may either reinforce or work against this tendency to underbid. In KL's example the bidding trade-off works against it, and with common values a random $L2$ coincides with equilibrium. With uniform, independent private values, a random $L2$ coincides with equilibrium, like a random $L1$. With non-uniform independent private values a random $L2$ may underbid, overbid, or coincide with equilibrium. In both of GHP's (2002) treatments, a random $L2$ tends to underbid.

For empirically plausible type distributions, with $L1$ more frequent than $L2$, these patterns of variation in bidding behavior allow a random level- k model to fit experimental data for common-value auctions significantly better than equilibrium or cursed equilibrium; and to fit GHP's data for non-uniform independent-private-value auctions significantly better than equilibrium and

cursed equilibrium.¹³ A level- k model has additional advantages over cursed equilibrium, in that it uses more general strategic principles to give a structural account of subjects' heterogeneous bidding behavior, with behavioral parameters linked to other bodies of evidence; and it has the potential to explain non-equilibrium bidding in some other independent-private-value auctions.¹⁴

The rest of the paper is organized as follows. In Section 2 we introduce MW's (1982) general model with interdependent values and affiliated signals and review the theories of equilibrium and cursed-equilibrium bidding. In Section 3 we discuss the specification of a level- k model for auctions and derive its implications for random and truthful type hierarchies. In Section 4 we compare equilibrium, cursed equilibrium, and the level- k model's implications with regard to explaining the curse and overbidding in independent-private-value auctions. We start with the two leading examples that were the basis of the auction experiments ER (2002, 2005) considered, the first-price auctions of KL (1986) and Garvin and Kagel (1994) and the second-price auctions of Avery and Kagel (1997; "AK"). We also analyze second-price auctions in KL's example (for which ER (2002) but not ER (2005) discuss KL's results). Further, since independent-private-value auctions are especially useful for separating cursed equilibrium from level- k decision rules, we analyze GHP's (2002) experimental design with discrete, slightly non-uniform values, in which level- k decision rules are separated from equilibrium. In Section 5 we reconsider the experimental data in the light of Section 4's analysis, comparing the models' abilities to account for the individual initial responses of inexperienced subjects, which allow the cleanest tests of our model, in settings like KL's, AK's, and GHP's. Section 6 is the conclusion.

2. Equilibrium and Cursed Equilibrium

In this section we review the theories of equilibrium and cursed-equilibrium bidding in first- and second-price auctions. We use MW's (1982, Section 3) general model with interdependent values and affiliated signals, which includes independent private values, pure common values, and intermediate cases in which bidders observe affiliated private signals that are informative about their interdependent values. Although ER's (2005) cursed equilibrium includes

¹³In common-value auctions a truthful LI tends to underbid relative to equilibrium or coincide with it in the examples, which makes it difficult for a truthful level- k model to reconcile the data with a plausible type distribution. However, our estimates for GHP's treatments also include a small but significant fraction of truthful LI subjects.

¹⁴Cursed equilibrium, by contrast, can accommodate subjects' heterogeneous bidding behavior only by estimating subject-specific cursedness parameters as in ER (2005, Section 4).

equilibrium as a special case, we begin with equilibrium and generalize to cursed equilibrium. Here and below, we assume risk-neutral, symmetric bidders and focus on symmetric equilibria.

2a. Milgrom and Weber's general model with interdependent values and affiliated signals

Milgrom and Weber's general model with interdependent values and affiliated signals has N bidders, indexed $i = 1, \dots, N$, bidding for a single, indivisible object. Bidder i observes a private signal X_i that is informative about his value of the object, with $X = (X_1, X_2, \dots, X_N)$. The vector $S = (S_1, S_2, \dots, S_M)$ includes additional random variables that may be informative about the value of the object. In general, bidder i 's value $V_i = u(S, X)$, and the variables in S and X are affiliated (positively associated) as defined in MW (1982, Assumption 5 and Appendix). This general model includes three leading special cases that are important in our analysis: the pure independent-private-value model, in which $M = 0$ and $V_i = u(X_i)$; the pure common-value model (used in KL (1986) and LP (1991)), in which $M = 1$ and $V_i = u(S)$; and an alternative common-value model (used in AK (1997)), in which $V_i = u(X) = \sum_{n=1}^N X_n$.

Because bidders are risk-neutral, if bidder i wins the auction and pays price p for the object his payoff is $V_i - p$. Without loss of generality, we focus on bidder 1 with value V_1 and signal X_1 . Let Y denote the highest estimate among other bidders' private signals X_2, X_3, \dots, X_N . Y has distribution function and density function, conditional on the realization, x , of X_1 , $F_Y(y | x)$ and $f_Y(y | x)$; in cases where the signals are independent, we suppress the conditioning and write $F_Y(y)$ and $f_Y(y)$. It is also useful to define two expected value functions, which as functions are the same for all i : the expected value conditional on winning $v(x, y) \equiv E[V_i | X_i = x, Y = y]$, and the unconditional expected value $r(x) = E[V_i | X_i = x]$.

2b. Equilibrium in first- and second-price auctions

Our equilibrium analysis closely follows MW's analysis of their general affiliated-signals and interdependent-values model; readers who are familiar with their analysis can skip ahead to Section 2.c's review of cursed equilibrium.

In equilibrium, bidders correctly predict and best respond to the distribution of other bidders' bids, taking into account the information to be revealed because in a symmetric equilibrium, the winner's estimate must be more favorable than others' estimates. In this

subsection we assume other bidders use their equilibrium bidding strategies, $a_*(x)$ in a first-price or $b_*(x)$ in a second-price auction, which are both increasing, with inverses $a_*^{-1}(a)$ and $b_*^{-1}(b)$.

In a first-price auction, bidder i 's optimal bidding strategy solves (for each x)

$$(1) \quad \max_a E[(V_i - a)1_{\{a_*(Y) < a\}} | X_i = x] = \int_{\underline{x}}^{a_*^{-1}(a)} (v(x, s) - a) f_Y(s | x) ds,$$

where $1_{\{\cdot\}}$ is the indicator function. Taking the partial derivative with respect to a yields a first-order differential equation that determines a as a function $a(x)$ of x , which characterizes the first-price equilibrium bidding strategy:¹⁵

$$(2) \quad a'(x) = (v(x, x) - a(x)) \frac{f_Y(x | x)}{F_Y(x | x)}.$$

Solving (2) for the equilibrium bidding strategy $a_*(x)$ and using the boundary condition $a_*(\underline{x}) = v(\underline{x}, \underline{x})$ to determine the constant of integration yields a general expression for the first-price equilibrium bidding strategy (MW (1982, p. 1107)):

$$(3) \quad a_*(x) = v(x, x) - \int_{\underline{x}}^x \exp\left(-\int_s^x \frac{f_Y(t | t)}{F_Y(t | t)} dt\right) d(v(s, s)).$$

$a_*(x)$ reflects both the value adjustment for the information revealed by winning, via $v(x, x)$, and the bidding trade-off, via the range of integration. The logic of value adjustment is that the bidder should bid according to the expected value given his own signal, conditional on *just* winning, which in equilibrium happens when his signal just exceeds the highest of the others' signals.

With independent private values, $v(x, x) = x$ and the functions $f_Y(y | x)$ and $F_Y(y | x)$ no longer depend on x , so the interior integral on the right-hand side of (3) reduces to $\frac{F_Y(s)}{F_Y(x)}$ and the first-price equilibrium bidding strategy becomes:

$$(4) \quad a_*(x) = x - \int_{\underline{x}}^x \frac{F_Y(s)}{F_Y(x)} ds \equiv E[Y | Y < X].$$

In a second-price auction, bidder i 's optimal bidding strategy solves (for each x):

¹⁵MW (1982, p. 1107-1108) show that the objective function in (1) is quasiconcave, so that the first-order conditions characterize the equilibrium strategies. MW's quasiconcavity argument breaks down for some of the optimization problems considered below, and level- k types' non-equilibrium beliefs can in general lead to boundary optima. In Section 4's examples the first-order conditions characterize the optimum except for the random $L2$ and truthful $L1$ types in AK's example, in which the objective function is linear and so either the upper or lower bound is optimal.

$$(5) \quad \max_b E[(V_i - b_*(Y)) \mathbb{1}_{\{b_*(Y) < b\}} | X_i] = \int_x^{b_*^{-1}(b)} (v(x, s) - v(s, s)) f_Y(s | x) ds .$$

Because $v(x, y)$ is increasing in x , $v(x, s) - v(s, s) > 0$ for all $s < x$ and $v(x, s) - v(s, s) < 0$ for all $s > x$. Thus the second-price equilibrium bidding strategy (MW (1982, pp. 1100-1101)) is:

$$(6) \quad b_*(x) = v(x, x) .$$

With independent private values, (6) becomes:

$$(7) \quad b_*(x) = x .$$

In a second-price auction a bidder's bid determines only when he wins, not what he pays, so the bidding trade-off is neutral and truthful bidding given correct value adjustment ensures that he wins if and only if it appears profitable, given his information. Comparing (3) to (4) and (6) to (7), the only differences between the common- and independent-private-value equilibrium bidding strategies are value adjustment and the affiliation of signals $f_Y(y | x)$. Although the independent-private-value equilibrium $b_*(x)$ is a weakly dominant strategy and the common-value equilibrium $b_*(x)$ is also truthful, it is not weakly dominant because optimal value adjustment depends on others' bidding strategies, as Section 3's level- k analysis shows more concretely.

Comparing (4) and (7) and recalling that in symmetric equilibrium a bidder wins if and only if he has the highest value yields the part of the revenue-equivalence theorem that is relevant to our analysis: With independent private values, in a first-price auction the winning bidder pays $E[Y | Y < X]$ while in a second-price auction he pays Y . Because $Y < X$ for the winner in either case, first- and second-price auctions yield the same expected revenue: random for second-price auctions, deterministic for first-price (Vickrey (1961)). By contrast, comparing (3) and (6) shows that with affiliated signals, a second-price auction yields at least as much expected revenue as a first-price auction (MW (1982, Theorem 15), McAfee and McMillan (1987, Sections V and X).

2c. Cursed equilibrium in first- and second-price auctions

Our cursed-equilibrium analysis follows ER's (2002, 2005) analysis; readers who are already familiar with it can skip ahead to Section 3's discussion of the level- k model.

In cursed equilibrium, as in equilibrium, bidders correctly predict and best respond to the distribution of others' bids. The only difference is that in cursed equilibrium bidders do not correctly perceive how others' bids depend on their signals. Instead they believe that with probability χ , ER's (2005) level of "cursedness," each other bidder bids the average of others' bids

over all signals rather than the bid his strategy specifies for his own signal. The parameter χ ranges from 0 to 1, and cursed equilibrium for a given χ is called " χ -cursed" equilibrium. $\chi = 0$ yields standard equilibrium and $\chi = 1$ yields "fully-cursed" equilibrium, in which bidders assume there is no relation between others' bids and signals, so that each takes the expected value of the object conditional on his own signal, ignoring the information revealed by winning.¹⁶

ER (2005, proof of Proposition 1, Proposition 5) simplify their analysis by showing that χ -cursed equilibrium is the same as equilibrium in a hypothetical " χ -virtual game," in which players believe that with probability χ others' bids are type-independent, in which case they learn nothing about the value of the object from winning. In the χ -virtual game, bidder i 's expected payoff from winning and paying price p when the value of the object is $u(S, X)$ is:

$$(8) \quad (1 - \chi)E[V_i | X_i = x, Y = x] + \chi E[V_i | X_i = x] = (1 - \chi)v(x, x) + \chi r(x).$$

The χ -cursed-equilibrium bidding strategy can then be obtained from the χ -virtual game in exactly the same way that the equilibrium bidding strategy was obtained from the original game.

With independent private values, $v(x, x) = r(x) = x$, the χ -virtual game reduces to the original game, and cursed equilibrium coincides with equilibrium. But with common values, $v(x, x)$ differs from $r(x)$, and cursed equilibrium differs from equilibrium. In this subsection we assume that other bidders use their χ -cursed-equilibrium bidding strategy, $a_\chi(x)$ in a first-price or $b_\chi(x)$ in a second-price auction, which are both increasing, with inverses $a_\chi^{-1}(a)$ and $b_\chi^{-1}(b)$.

In a first-price auction bidder i 's optimal bidding strategy solves (for each x):

$$(9) \quad \max_a \int_x^{a_\chi^{-1}(a)} ((1 - \chi)v(x, s) + \chi r(x) - a) f_Y(s | x) ds.$$

Taking the partial derivative yields a first-order differential equation that determines a as a function $a(x)$, which characterizes the first-price χ -cursed-equilibrium bidding strategy:

$$(10) \quad a'(x) = ((1 - \chi)v(x, x) + \chi r(x)) - a(x) \frac{f_Y(x | x)}{F_Y(x | x)}.$$

Solving (10) yields:

¹⁶The implicit assumption that a player thinks he is more sophisticated than other players is often seen in other forms, for which it has considerable experimental support; see for example Weizsäcker (2003). As ER (2005, footnote 6) note, cursed equilibrium allows certain kinds of differences in beliefs about others' type-contingent strategies.

$$(11) \quad a_x(x) = [(1 - \chi)v(x, x) + \chi r(x)] - \int_x^x \exp\left(-\int_s^x \frac{f_Y(t|t)}{F_Y(t|t)} dt\right) d[(1 - \chi)v(s, s) + \chi r(s)].$$

Like the first-price equilibrium bidding strategy $a_*(x)$, $a_x(x)$ reflects both value adjustment and the bidding trade-off. Cursed equilibrium differs from equilibrium only in underestimating the correct value adjustment to an extent determined by χ .

Given a cursed-equilibrium bidder's value estimate and its anticipation of others' estimates, he responds to the bidding trade-off just as an equilibrium bidder would. The effect of its cursedness is determined by the difference between the unconditional expected value $r(x)$ and the expected value conditional on winning $v(x, x)$. Normally $r(x) > v(x, x)$, so that cursed-equilibrium bidders overbid, relative to equilibrium, as in KL's example (Section 4). But there are some cases in which $v(x, x) > r(x)$ for some values of x , so that some (in extreme cases, nearly all) cursed-equilibrium bidders underbid, as in AK's example (Section 4; ER (2005, p. 22)).

In a second-price auction, bidder i 's optimal bidding strategy solves (for each x):

$$(12) \quad \max_b \int_x^{b_x^{-1}(b)} ((1 - \chi)v(x, s) + \chi r(x) - (1 - \chi)v(s, s) - \chi r(s)) f_Y(s | x) ds,$$

which (following the same reasoning as for equilibrium, because both $v(x, y)$ and $r(x)$ are increasing in x) yields the second-price χ -cursed-equilibrium bidding strategy:

$$(13) \quad b_x(x) = (1 - \chi)v(x, x) + \chi r(x).$$

Like the second-price equilibrium bidding strategy $b_*(x)$, $b_x(x)$ reflects only the value adjustment for the information revealed by winning, which it underestimates just as in a first-price auction.

Assuming χ is the same in first- and second-price auctions, the revenue implications are qualitatively the same for cursed equilibrium as for equilibrium: With independent private values, first- and second-price auctions yield the same expected revenue; but with affiliated signals, a second-price auction yields at least as much expected revenue as a first-price auction.

3. The Level- k Model

3a. The model

In this section we generalize the level- k model to common- and independent-private-value auctions. As explained in the Introduction, the level- k model allows behavior to be heterogeneous across bidders, but assumes that each bidder's behavior is drawn from a common distribution over

a hierarchy of types, in which $L1$ best responds to a nonstrategic anchoring type $L0$, $L2$ best responds to $L1$, etc. In this section we derive types' implications in general; and in Section 4 we specialize them to the examples used in experiments and discuss aggregate implications. We consider two alternative specifications of $L0$: a *random* $L0$ that bids uniformly randomly, independent of its own private signal, over the range determined by the range of its signal and the value function $V_i = u(S, X)$; and a *truthful* $L0$ that bids the value its own signal suggests, taken by itself. We assume that each player follows type $L0$, $L1$, or $L2$ (footnote 9), random or truthful. (Recall that we call the $L1$ and $L2$ types associated with a random $L0$, random $L1$ and $L2$ types, with analogous terms for the truthful $L1$ and $L2$ types associated with a truthful $L0$; and that random or truthful $L1$ and $L2$ types need not be random or truthful themselves.)

3b. Random $L1$ and $L2$ bidding strategies in first- and second-price auctions

A random $L1$ bidder assumes other bidders are random $L0$, hence with bids independently and identically distributed (henceforth "i.i.d.") uniformly over the range $[\underline{z}, \bar{z}]$ determined by the range of its private signal and the value function $V_i = u(S, X)$. A random $L1$ therefore believes that winning conveys no information about the value of the object, even with common values and affiliated signals. Its optimal bid is determined by its own signal; the price it pays if it wins; and its beliefs about the highest bid among the others, Z_1 , described by the distribution function

$F_{Z_1}(z) = \left(\frac{z - \underline{z}}{\bar{z} - \underline{z}}\right)^{N-1}$ and density $f_{Z_1}(z) = (N-1) \left(\frac{z - \underline{z}}{\bar{z} - \underline{z}}\right)^{N-2} \frac{1}{\bar{z} - \underline{z}}$. Note that these do not depend on the bidder's own signal X_1 , which is uninformative about Z_1 ; or the distribution of others' signals.

In a first-price auction a random $L1$ bidder i 's optimal bidding strategy solves (for each x):

$$(14) \quad \max_a E \left[(V_i - a) \mathbf{1}_{\{Z_1 < a\}} \mid X_i \right] \equiv \int_{\underline{z}}^a (r(x) - a) f_{Z_1}(s) ds \equiv (r(x) - a) F_{Z_1}(a).$$

A random $L1$'s first-price bidding strategy, $a_1^r(x)$, is characterized by the first-order condition:

$$(15) \quad (r(x) - a) f_{Z_1}(a) - F_{Z_1}(a) = 0.$$

This problem and first-order condition differ from those for first-price equilibrium in (1) and (2) in two ways: $r(x)$ replaces $v(x, x)$, and the integral in (14) and density and distribution function in (15) refer to a random $L1$'s beliefs about the highest of $L0$ others' bids Z_1 , rather than the highest of others' signals Y that determines the highest others' bid in a symmetric equilibrium. The first

difference reflects the fact that because a random LI 's bids are independent of its private signal, a random LI ignores the information revealed by winning. Given the normal tendency for $r(x) > v(x, x)$, this tends to make a random LI overbid relative to equilibrium, just as a fully-cursed equilibrium bidder does. The second difference reflects a random LI 's use of its non-equilibrium beliefs to evaluate the bidding trade-off between a higher bid's cost and increased probability of winning. Comparing (10) with $\chi = 1$ and (15) shows that this difference can either raise or lower its first-price bidding strategy relative to the fully-cursed equilibrium bidding strategy. Combining the two tendencies, a random LI seems more likely than not to overbid relative to equilibrium.

In a second-price auction, a random LI bidder i 's optimal bidding strategy solves:

$$(16) \quad \max_b E[(V_i - Z_1)1_{\{Z_1 < b\}} | X_i] = \int_{\underline{z}}^b (r(x) - s) f_{Z_1}(s) ds$$

A random LI 's second-price bidding strategy, $b_1^r(x)$, is characterized by the first-order condition:

$$(17) \quad (r(x) - b) f_{Z_1}(b) = 0 \text{ or, solving for } b, b_1^r(x) = r(x).$$

This problem and first-order condition differ from those for second-price equilibrium in (5) and (6) in that $r(x)$ replaces $v(x, x)$ and in the use of a random LI 's non-equilibrium beliefs. But given a random LI 's "cursed" value adjustment, truthful bidding is optimal just as it is in an equilibrium analysis.¹⁷ This important insight from an equilibrium analysis remains valid here and below, even though the truthful equilibrium bidding strategy in (6) is not weakly dominant and random LI beliefs differ from equilibrium beliefs, because a bidder's bid in a second-price auction still determines only when he wins, not what he pays; and truthful bidding, given correct value adjustment taking others' anticipated bidding strategies into account, still ensures that he wins when it appears profitable, given his beliefs. A random LI 's bidding strategy therefore coincides with the second-price fully-cursed equilibrium bidding strategy in (13) with $\chi = 1$, so that it has the same tendency to overbid in common-value auctions. But it coincides with equilibrium in second-price independent-private-value auctions, where like other level- k types with $k > 0$, which all best respond to beliefs, it follows the weakly dominant strategy.

Unlike a random LI , a random $L2$ adjusts its value estimate for the information revealed by winning, because a random LI 's bidding strategy is an increasing function of its private signal

¹⁷Fully-cursed equilibrium and random LI are readily comparable because both are determined by the unconditional expected value $r(x)$ instead of the value conditional on just winning $v(x, x)$, and so differ only in their beliefs. Even so, in first-price auctions random LI and fully-cursed equilibrium are not directly comparable, because random LI 's and equilibrium beliefs can differ considerably, depending on the specific distribution of the signals.

in either kind of auction.¹⁸ We derive the optimal bids more generally, because the results will determine the truthful $L1$ and $L2$ bidding strategies as well as the random $L2$ bidding strategy.

Suppose that in a first-price auction, a level- k bidder (random or truthful) expects others to bid according to the monotonically increasing bidding strategy $a_{k-1}(x)$, with inverse $a_{k-1}^{-1}(a)$. The bidder's optimal bidding strategy with value V_i and signal X_i then solves (for each x):

$$(18) \quad \max_a E[(V_i - a) \mathbb{1}_{\{a_{k-1}(Y) < a\}} | X_i] = E[(V_i - a) \mathbb{1}_{\{Y < a_{k-1}^{-1}(a)\}} | X_i] = \int_x^{a_{k-1}^{-1}(a)} (v(x, s) - a) f_Y(s | x) ds.$$

Taking the partial derivative with respect to a , the first-order condition can be written:

$$(19) \quad (v(x, a_{k-1}^{-1}(a)) - a) f_Y(a_{k-1}^{-1}(a) | x) \frac{\partial a_{k-1}^{-1}(a)}{\partial a} - F_Y(a_{k-1}^{-1}(a) | x) = 0.$$

With independent private values $v(x, x) = x$ and the functions $f_Y(y | x)$ and $F_Y(y | x)$ no longer depend on x , so that (19) reduces to:

$$(20) \quad (x - a) f_Y(a_{k-1}^{-1}(a)) \frac{\partial a_{k-1}^{-1}(a)}{\partial a} - F_Y(a_{k-1}^{-1}(a)) = 0 \quad \text{or} \quad (x - a) = \frac{F_Y(a_{k-1}^{-1}(a))}{f_Y(a_{k-1}^{-1}(a))} \bigg/ \frac{\partial a_{k-1}^{-1}(a)}{\partial a}.$$

Now suppose that in a second-price auction, a level- k bidder expects others to follow the monotonic bidding strategy $b_{k-1}(x)$, with inverse $b_{k-1}^{-1}(b)$. The bidder's optimal bidding strategy with value V_i and signal X_i then solves (for each x):

$$(21) \quad \max_b E[(V_i - b_{k-1}(Y)) \mathbb{1}_{\{b_{k-1}(Y) < b\}} | X_i] = \int_x^{b_{k-1}^{-1}(b)} (v(x, s) - b_{k-1}(s)) f_Y(s | x) ds.$$

Taking the partial derivative with respect to b , the first-order condition can be written:

$$(22) \quad (v(x, b_{k-1}^{-1}(b)) - b) f_Y(b_{k-1}^{-1}(b) | x) \frac{\partial b_{k-1}^{-1}(b)}{\partial b} = 0 \quad \text{or} \quad v(x, b_{k-1}^{-1}(b)) - b = 0.$$

With independent private values (22) reduces to the weakly dominant strategy in (7).

Comparing the second-price level- k bidding strategy from (22) with the second-price equilibrium bidding strategy from (6) isolates the effects of value adjustment. The logic of value adjustment is the same for both: Each bids according to the expected value given its own signal, conditional on *just* winning. The only difference is that a level- k bidder's beliefs do not anticipate winning if and only if it has the highest signal, as a (symmetric) equilibrium bidder's do. A level- k

¹⁸This is easily verified from (15) for first-price auctions and (17) for second-price auctions.

bidder believes it wins if and only if it bids at least $b_{k-1}(Y)$, which depending on others' anticipated bidding strategy may be more or less stringent than having the highest signal.

Value adjustment tends to make bidders' bids strategic substitutes. Suppose that a level- k bidder believes others' bids are higher than in equilibrium, so winning means others' signals are (stochastically) lower than it would mean in equilibrium. Comparing (22) and (6) and noting that $v(x, y)$ is increasing in y (MW (1982, Theorems 2-5)), this belief lowers his value conditioned on winning, making the curse seem worse and lowering his optimal bid, other things equal.

Comparing the first-price level- k bidding strategy determined by (19) with the first-price equilibrium bidding strategy determined by (2) reveals that both involve exactly the same kind of value adjustment as in the second-price bidding strategies. In first-price auctions, however, the value adjustment interacts with the bidding trade-off. We now investigate this interaction in more detail, in preparation for Section 4's analysis of examples.

First, isolate the bidding trade-off by considering the level- k bidding strategy with independent private values determined by (20), which balances the marginal benefit of increasing its bid (the value minus the bid times the increased probability of winning) against the marginal cost (the higher bid times the probability of winning).¹⁹

Now write $a_{k-1}(x, q)$ as a function of a parameter q , where increasing q shifts $a_{k-1}(x, q)$ up for all x , and so shifts $a_{k-1}^{-1}(a, q)$ down for all a . With independent private values, (18) becomes:

$$(23) \quad \max_a (x - a)F(a_{k-1}^{-1}(a, q)).$$

The first-price level- k bidding strategy is then:

$$(24) \quad \arg \max_a (x - a)F(a_{k-1}^{-1}(a, q)) \equiv \arg \max_a \{\log(x - a) + \log[F(a_{k-1}^{-1}(a, q))]\},$$

where \log is the natural logarithm and $F(\cdot) > 0$ near the optimum. Because q enters only the second term of the right-hand maximand, the optimal bid a is everywhere increasing in q iff:

$$(25) \quad \partial^2 \log[F(a_{k-1}^{-1}(a, q))] / \partial a \partial q \equiv \frac{\partial \left[\frac{f(a_{k-1}^{-1}(a, q))}{F(a_{k-1}^{-1}(a, q))} \frac{\partial a_{k-1}^{-1}(a, q)}{\partial a} \right]}{\partial q} \geq 0 \text{ for all } a \text{ and } q,$$

or equivalently (given that all terms are positive) iff:

¹⁹Private values are enough to isolate the bidding trade-off; we add independence for convenience.

$$(26) \quad \frac{\partial \left[\frac{F(a_{k-1}^{-1}(a, q))}{f(a_{k-1}^{-1}(a, q))} \right] / \frac{\partial a_{k-1}^{-1}(a, q)}{\partial a}}{\partial q} \leq 0 \text{ for all } a, x, \text{ and } q,$$

with an analogous condition for the optimal bid to be everywhere decreasing in q . The numerator in the square brackets on the left-hand side of (26) is decreasing in q for most well-behaved distributions; but the denominator is also likely to be decreasing in q . Thus, neither the condition in (26) nor its converse are satisfied for all plausible specifications of $F(\cdot)$ and $a_{k-1}(\cdot, q)$; and the condition is likely to be satisfied for some values of a but not others: In general the bidding trade-off may make bidders' bids strategic complements, strategic substitutes, or a mixture of both.

To see what determines the effect of the bidding trade-off more clearly, assume that

$$a_{k-1}(x) \equiv \gamma x + \delta \text{ with } \gamma > 0 \text{ as in KL's and AK's examples, so that } a_{k-1}^{-1}(a) \equiv \frac{a - \delta}{\gamma} \text{ and}$$

$$\frac{\partial a_{k-1}^{-1}(a)}{\partial a} = \frac{1}{\gamma}. \text{ Given this and (26), increasing } \gamma \text{ increases the optimal bid } a \text{ iff it}$$

decreases $\gamma F\left(\frac{a - \delta}{\gamma}\right) / f\left(\frac{a - \delta}{\gamma}\right)$, which will be the case, given $F(x) = 0$, iff $\frac{F_Y(y)}{f_Y(y)}$ is convex in y .

But an increase in δ will increase the optimal bid a iff it decreases $F\left(\frac{a - \delta}{\gamma}\right) / f\left(\frac{a - \delta}{\gamma}\right)$, which is

the case for most well-behaved distributions. Thus, upward shifts in the slope of others' anticipated bidding strategy tend to make bidders' bids strategic complements (respectively substitutes) iff $\frac{F_Y(y)}{f_Y(y)}$ is convex (concave) in y , while upward shifts in the level tend to make

bidders' bids strategic complements in either case. With an unconditionally uniform signal distribution, as in most independent-private-value examples, $\frac{F_Y(y)}{f_Y(y)}$ is linear, and the bidding

trade-off is neutral with respect to shifts in the slope.

The ambiguity just demonstrated for the independent-private-value case plainly persists in the common-value case, where the bidding trade-off interacts with value adjustment: In general, comparing (19) with (2), the first-price level- k bidding strategy can be either higher than the first-price equilibrium bidding strategy, lower, or a mixture of both.

Now consider how a random $L2$'s first-price bidding strategy, $a_2^r(x)$, is determined by (19) with $a_1^{r^{-1}}(a)$ replacing $a_{k-1}^{-1}(a)$, hence by:

$$(27) \quad (v(x, a_1^{r^{-1}}(a)) - a) f_Y(a_1^{r^{-1}}(a) | x) \frac{\partial a_1^{r^{-1}}(a)}{\partial a} - F_Y(a_1^{r^{-1}}(a) | x) = 0.$$

In a first-price auction, a random $L2$ bidder, like a random $L1$ bidder, deviates from equilibrium both in value adjustment and in using its non-equilibrium beliefs to evaluate the bidding trade-off. A random $L2$'s value adjustment reflects the same logic as an equilibrium bidder's, but its beliefs generally lead to a different adjustment. To the extent that a random $L1$ overbids relative to equilibrium, because a random $L2$ believes that to win it must bid higher than all others' random $L1$ bids, not just higher than their equilibrium bids, given the strategic substitutability of value adjustment a random $L2$ believes that the curse is more severe than in equilibrium, and so tends to underbid. In general the bidding trade-off may tend to raise or lower a random $L2$'s bids relative to equilibrium or cursed equilibrium. On balance, a random $L2$ seems more likely to underbid.

A random $L2$'s second-price bidding strategy, $b_2^r(x)$, is determined by (22) with $b_1^{r^{-1}}(b)$ replacing $b_{k-1}^{-1}(b)$, hence by:

$$(28) \quad (v(x, b_1^{r^{-1}}(b)) - b) f_Y(b_1^{r^{-1}}(b) | x) \frac{\partial b_1^{r^{-1}}(b)}{\partial b} = 0 \text{ or } b = v(x, b_1^{r^{-1}}(b)).$$

The second-price random $L2$ bidding strategy is again truthful, but to the extent that a random $L1$ overbids relative to equilibrium, the strategic substitutability of value adjustment makes a random $L2$ underbid because it believes the curse is more severe than in equilibrium.

3c. Truthful $L1$ and $L2$ bidding strategies in first- and second-price auctions

A truthful $L1$ bidder's bid is a best response to a truthful $L0$, and thus assumes that others follow the monotonic bidding strategy $a_0^t(x) \equiv r(x) = E[V_i | X_i = x]$, with inverse $a_0^{t^{-1}}(a) \equiv r^{-1}(a)$.

In a first-price auction, a truthful $L1$ bidder's optimal bidding strategy, $a_1^t(x)$, solves a problem (for each x) that is a special case of the general first-price monotonic problem (18).

$a_1^t(x)$ is then determined by the first-order condition (19) with $a_0^{t^{-1}}(a) \equiv r^{-1}(a)$ (because $a_0^t(x) \equiv r(x)$) replacing $a_{k-1}^{-1}(a)$:

$$(29) \quad (v(x, r^{-1}(a)) - a) f_Y(r^{-1}(a) | x) \frac{\partial r^{-1}(a)}{\partial a} - F_Y(r^{-1}(a) | x) = 0.$$

Thus, in a first-price auction, a truthful *LI* bidder deviates from equilibrium in its use of its non-equilibrium beliefs to evaluate the bidding trade-off, like a random *LI*; but its different beliefs imply a different value adjustment.²⁰ A truthful *LO* overbids relative to the first-price equilibrium bidding strategy, because it neither adjusts for the curse nor shades its bids. Hence a truthful *LI*, which believes that to win it must bid higher than all others' truthful bids, not just higher than their first-price equilibrium bids, believes that the curse is even more severe than in equilibrium. Thus the strategic substitutability of value adjustment makes a truthful *LI* tend to underbid. The bidding trade-off, by contrast, may in general either raise or lower a truthful *LI*'s bids relative to equilibrium. On balance, a truthful *LI* seems more likely to underbid.

In a second-price auction, a truthful *LI* bidder's optimal bidding strategy, $b_1^t(x)$, solves a special case of the general monotonic problem (21) (for each x). The bidder's optimal second-price bidding strategy, $b_1^t(x)$, is then determined by (22) with $b_0^{-1}(b) \equiv r^{-1}(b)$ replacing $b_{k-1}^{-1}(b)$:

$$(30) \quad (v(x, r^{-1}(b)) - b) f_Y(r^{-1}(b) | x) \frac{\partial r^{-1}(b)}{\partial b} = 0 \text{ or } b = v(x, r^{-1}(b)).$$

Thus, bidding is truthful as in the previous second-price analyses. A truthful *LO* normally overbids relative to second-price equilibrium because it does not adjust for the curse, hence the strategic substitutability of value adjustment normally makes a truthful *LI* underbid. In a common-value second-price auction, a truthful *LI*'s bidding strategy is identical to a random *L2*'s, because a random *LI* bids the expected value of the item based on its own signal, just as a truthful *LO* does.

In a first-price auction, a truthful *L2* bidder expects other bidders to bid according to the monotonic bidding strategy $a_1^t(x)$, with inverse $a_1^{t-1}(a)$. Its optimal first-price bidding strategy, $a_2^t(x)$, is then determined by problem (18) with $a_1^{t-1}(a)$ replacing $a_{k-1}^{-1}(a)$:

$$(31) \quad (v(x, a_1^{t-1}(a)) - a) f_Y(a_1^{t-1}(a) | x) \frac{\partial a_1^{t-1}(a)}{\partial a} - F_Y(a_1^{t-1}(a) | x) = 0.$$

Thus, to the extent that a truthful *LI* underbids, the strategic substitutability of value adjustment tends to make a truthful *L2* overbid. But the bidding trade-off can again either raise or lower its bids relative to equilibrium. On balance, a truthful *L2* seems likely to overbid.

²⁰Because truthful types' bidding strategies are determined by $v(x,y)$, like equilibrium strategies, they are more readily compared to equilibrium than to cursed-equilibrium strategies, which are influenced by $r(x)$ as well as $v(x,y)$.

In a second-price auction, a truthful $L2$ bidder expects other bidders to bid according to the monotonic bidding strategy $b_1^t(x)$, with inverse $b_1^{t-1}(b)$. Its optimal second-price bidding strategy, $b_2^t(x)$, is again determined by (22), now with $b_1^{t-1}(b)$ replacing $b_{k-1}^{-1}(b)$:

$$(32) \quad (v(x, b_1^{t-1}(b)) - b) f_Y(b_1^{t-1}(b) | x) \frac{\partial b_1^{t-1}(b)}{\partial b} = 0 \text{ or } b = v(x, b_1^{t-1}(b)).$$

To the extent that a truthful $L1$ underbids, the strategic substitutability of value adjustment again tends to make a truthful $L2$ overbid.

4. Can a Level- k Model Explain the Curse and Other Kinds of Overbidding?

The auction experiments whose data we analyze are based on two leading common-value examples, which differ mainly in the form of the value function $V_i = u(S, X)$ (Section 2); and one independent-private-value example. In this section we review the implications of equilibrium, cursed equilibrium, and the level- k model in general and discuss them in detail in the examples, to assess the level- k model's potential to improve upon cursed-equilibrium explanations of the curse in common-value auctions and other kinds of overbidding in independent-private-value auctions.

4a. Kagel and Levin's; Avery and Kagel's; and Goeree, Holt, and Palfrey's examples

In the first example, used in KL's analyses of first-price auctions and in LP's follow-up experiments, $N \geq 3$, $V_i = u(S, X) = S$, S is uniformly distributed on a subset of the real line $[\underline{s}, \bar{s}]$, and $X|S$ is conditionally uniformly i.i.d. on the interval $[s - \frac{a}{2}, s + \frac{a}{2}]$ with dispersion $a > 0$, with minor adjustments due to truncation near \underline{s} or \bar{s} . The density, distribution function, and expected value of $X|S$ are: $f_{X|S} = \frac{1}{a}$, $F_{X|S} = \frac{x-s}{a} + \frac{1}{2}$, and $E[X | S] = s$. Thus $r(x) \equiv E[S | X = x] = x$.

Standard calculations show that:

$$(33) \quad v(x, y) = \begin{cases} x - \frac{a}{2} + \frac{a}{N} - \frac{x-y}{N}, & x - a \leq y \leq x \\ y - \frac{a}{2} + \frac{a}{N} - \frac{\left(\frac{y-x}{a}\right)^{N-1}}{\left[1 - \left(\frac{y-x}{a}\right)^{N-1}\right]} \left(\frac{N-1}{N}\right) (x + a - y), & x < y \leq x + a \end{cases}.$$

Thus $v(x, x) = x - \frac{a}{2} + \frac{a}{N} \leq r(x) = x$, with strict inequality for $N > 2$, and cursed-equilibrium bidders overbid relative to equilibrium or coincide with it for any χ or x .

In the second example, used in AK's analysis of second-price auctions,

$V_i = u(S, X) = \sum_{i=1}^N X_i$, and X_i is i.i.d. uniformly distributed on the interval $[\underline{x}, \bar{x}]$. Thus, in general,

$$r(x) \equiv E\left[\sum_{k=1}^N X_k \mid X_i = x\right] = x + (N-1)\frac{\bar{x} + x}{2}, \quad v(x, y) = x + y\frac{N}{2} + \frac{(N-2)}{2}\underline{x}, \text{ and}$$

$$v(x, x) = x + x\frac{N}{2} + \frac{N-2}{2}\underline{x} > (<) r(x) \text{ if and only if } x > (<) \frac{(N-1)\bar{x} + x}{N}, \text{ so that } v(x, x) > r(x) \text{ for}$$

bidders with high signals and $v(x, x) < r(x)$ for bidders with low signals: Cursed-equilibrium bidders underbid relative to equilibrium for high signals (because they implicitly assume that others' signals take their average values, when their own signal makes others' more likely to be high) and overbid for low signals.²¹ When $N = 2$ and $[\underline{x}, \bar{x}] = [1, 4]$, as in AK's experiments,

$$r(x) = x + \frac{5}{2} \text{ and } v(x, x) = 2x, \text{ so that } r(x) < v(x, x) \text{ when } x > \frac{5}{2} \text{ and } r(x) > v(x, x) \text{ when } x < \frac{5}{2}.$$

In the third example, used in GHP's (2002) analysis of first-price independent-private-value auctions, $N = 2$, $V_i = u(S, X) = X_i$, and there are two treatments, each with bids restricted to integer values and discrete, slightly non-uniform (because of spacing) values—equal probabilities on $\{0, 2, 4, 6, 8, 11\}$ in a low-value treatment and on $\{0, 3, 5, 7, 9, 12\}$ in a high-value treatment.

4b. Equilibrium and cursed equilibrium versus level- k models in second-price auctions

We now describe and discuss the relationships among equilibrium, cursed-equilibrium, and random and truthful $L1$ and $L2$ bidding strategies in these examples. Table 1 summarizes the conclusions of Section 2's and Section 3's general analyses and records the models' implications in the examples.

First consider second-price auctions with independent private values. Here random and truthful $L1$ and $L2$ all bid truthfully, just as they do in equilibrium and cursed equilibrium, because level- k types follow weakly dominant strategies when they exist.²² Thus neither kind of

²¹This corrects a typographical error in ER (2005, p. 1642), where they say that bidders with high signals overbid relative to equilibrium while those with low signals underbid.

²²In this case a truthful $L0$ (but not a random $L0$) also coincides with equilibrium when a player's signal reveals the actual value with certainty.

level- k model can improve upon an equilibrium explanation of overbidding in second-price auctions with independent private values.

Now consider second-price auctions with common values. In this case, a random $L1$ bids the value its own signal suggests, taken by itself, so its bidding strategy coincides with fully-cursed equilibrium and it has the same general tendency to overbid relative to equilibrium (Section 3b).²³ In KL's example, a random $L1$ coincides with equilibrium for $N = 2$ and, like fully-cursed equilibrium, overbids relative to equilibrium for $N > 2$, to an extent that increases with N and the dispersion a of its signal. In AK's example, a random $L1$ with a low (high) signal overbids (underbids), again like fully-cursed equilibrium.

To the extent that a random $L1$ overbids relative to equilibrium in a second-price common-value auction, the strategic substitutability of value adjustment make a random $L2$ underbid. In KL's example, a random $L2$, like a fully-cursed equilibrium bidder, coincides with equilibrium for $N = 2$ and underbids for $N > 2$, to an extent that decreases with N and increases with a . In AK's example, a random $L2$ with a low (high) signal matches the bid of a random $L1$ with the lowest (highest) possible signal (with only weak strategic substitutability because the solutions are on the boundary).

In a second-price auction with common values, a truthful $L0$ generally overbids relative to equilibrium, because it does not adjust its value for the curse. Given this, the strategic substitutability of value adjustment makes a truthful $L1$ underbid relative to equilibrium or, a fortiori, cursed equilibrium. In KL's example, because a random $L1$ bids the value its own signal suggests, like a truthful $L0$, a truthful $L1$ coincides with a random $L2$, and so underbids relative to equilibrium, to an extent that decreases with N and increases with a . In AK's example with $N = 2$, a truthful $L1$ coincides with a random $L2$ (Section 3c), and therefore a truthful $L1$ with a low (high) signal matches the bid of a truthful $L0$ with the lowest (highest) possible signal (with only weak strategic substitutability because the solutions are on the boundary).

Given that a truthful $L1$ underbids, a truthful $L2$ overbids. We have not derived a closed-form solution for a truthful $L2$ in KL's or AK's example. Computations confirm that it overbids in

²³Recall that the effect of cursedness is determined by the difference between the unconditional expected value of the item $r(x)$ and the expected value of the item conditional on just winning $v(x,x)$ (Section 2c). Normally $r(x) > v(x,x)$, so that cursed-equilibrium bidders overbid, relative to equilibrium; but there are cases in which $v(x,x) > r(x)$ for some x , so that some cursed-equilibrium bidders underbid.

KL's example by even more than a fully cursed-equilibrium bidder, to an extent that increases with N ; and that in AK's example with $N = 2$ it overbids for some values and underbids for others.

To sum up for second-price auctions, with independent private values level- k types of either kind coincide with equilibrium and cursed equilibrium, and so cannot improve upon their explanations. With common values a random $L1$ coincides with fully-cursed equilibrium, and so has a general tendency to overbid relative to a less-than-fully-cursed-equilibrium (or equilibrium) bidder; and a random $L2$ has a general tendency to underbid relative even to equilibrium. Thus a random level- k model has the potential to improve upon χ -cursed-equilibrium explanations of the curse in second-price common-value auctions; but this depends on whether its mixture of $L1$ and $L2$ bidders gives a better account of subjects' heterogeneous bidding behavior, with empirically plausible type frequencies, than χ -cursed-equilibrium with a mixture of cursed-equilibrium bidders with different χ "types" with a comparable number of parameters.²⁴

By contrast, a truthful $L1$ has a general tendency to underbid; and a truthful $L2$ to overbid, sometimes by even more than a fully-cursed-equilibrium bidder. Given the prevalence of $L1$ over $L2$ types in experimental subject populations (see Costa-Gomes and Crawford (2005) and the papers mentioned there), this makes it difficult to reconcile a truthful level- k model with the frequency with which bidders fall into the curse.

4c. Equilibrium and cursed equilibrium versus level- k models in first-price auctions

We now discuss the relationships among equilibrium, cursed-equilibrium, and random and truthful $L1$ and $L2$ bidding strategies for first-price auctions. Here, with independent private values, value adjustment is irrelevant, and in general the bidding trade-off may make a random or truthful $L1$ or $L2$ either underbid or overbid.

To date, most independent-private-value experiments have used values uniformly i.i.d. on $[\underline{x}, \bar{x}]$. In this case the equilibrium bidding strategy $a_*(x) = \frac{n-1}{n}(x - \underline{x}) + \underline{x}$ is a best response to any beliefs derived from others' bidding strategies $a(x - \underline{x}) + \underline{x}$, as long as $0 < a \leq 1$. A random or truthful $L1$, and therefore a random or truthful $L2$, then coincides with equilibrium; and the level- k model cannot improve upon an equilibrium explanation of overbidding. But for more general value distributions a level- k model may be able to explain non-equilibrium bidding with

²⁴In KL's example, for instance, random $L2$'s or truthful $L1$'s below-equilibrium bids could only be duplicated by the deterministic part of an econometric specification of χ -cursed equilibrium if $\chi < 0$; and truthful $L2$'s above-signal bids could only be duplicated if $\chi > 1$.

independent private values. In GHP (2002), a random $L1$ or $L2$ coincides with equilibrium except for the highest valuation in the high-value treatment, where a random $L1$ overbids and a random $L2$ underbids (Appendix).²⁵ By contrast, a truthful $L1$ underbids in the low-value treatment and overbids in the high-value treatment, and a truthful $L2$ underbids for both. Random and truthful specifications do not coincide here, even though both the ex ante value distribution and a random $L0$ are uniformly distributed over the possible values, because the possible values are not evenly spaced and the allowed integer bids include some between the values.

Now consider first-price common-value auctions.²⁶ Value adjustment follows the same principle as in second-price auctions, and tends to make a random $L1$ overbid relative to equilibrium because it does not adjust for the curse. But now the bidding trade-off may either reinforce or offset this tendency. In general a random $L1$ may underbid or overbid; but on balance it seems more likely to overbid. In KL's example when $N = 2$, equilibrium and fully-cursed equilibrium coincide and a random $L1$ bids slightly lower than but approximately coincides with them.²⁷ When $N > 2$, a random $L1$ (approximately) coincides with a fully-cursed equilibrium bidder; but both overbid relative to equilibrium, by an amount that increases with N and the dispersion parameter a .

To the extent that a random $L1$ overbids in a first-price common-value auction, value adjustment makes a random $L2$ tend to underbid; but the bidding trade-off may either reinforce or offset this tendency. On balance it seems more likely to underbid. In KL's example when $N = 2$, a random $L2$ (approximately) coincides with equilibrium and fully-cursed equilibrium. When $N > 2$, a random $L2$ (approximately) coincides with equilibrium, but it underbids relative to fully-cursed equilibrium, by an amount that increases with N and a .²⁸

²⁵Although random $L1$ approximately coincides with equilibrium here, their different beliefs imply different costs of deviation, which plays an important role in the econometric analysis. In Palfrey's (1985) and Chen and Plott's (1998) designs, a level- k model also deviates from equilibrium, in different ways; we do not consider their designs here because the data from those experiments are not available.

²⁶Although in independent-private-value auctions, random types are equivalent to the analogous truthful types when the distribution of private signals is unconditionally uniform; in common-value auctions random and truthful types are never equivalent, because they differ in value adjustment.

²⁷"Approximately coincides" means that the bidding strategies differ only by the exponential part of KL's example's first-price equilibrium bidding strategy, which is positive but negligible for all x not very close to \underline{x} ; KL and all other analysts have ignored this exponential part and we will follow them in this from now on, for cursed equilibrium as well as equilibrium. Our solution for KL's example differs from KL's, LP's (1991), and ER's, in that they have $a/(N+1)$ in the third term in place of our a/N (Appendix). We believe that our version is correct, but the discrepancy makes little difference because the exponential term is negligible.

²⁸In this example the bidding trade-off tends to make bids strategic complements because it increases the intercept δ but not the slope γ (Section 3b), which offsets the strategic substitutability of value adjustment.

In first-price common-value auctions, a truthful $L0$ overbids relative to equilibrium, more than in a second-price auction, other things equal, because it neither adjusts its value for the curse nor shades its bid as the bidding trade-off requires. Because a truthful $L0$ overbids, with common values the strategic substitutability of value adjustment tends to make a truthful $L1$ underbid (Section 3.c); but the bidding trade-off may either reinforce or offset this tendency. In general a truthful $L1$ may either underbid or overbid. On balance, it seems more likely to underbid. In KL's example value adjustment and the bidding trade-off offset each other, and a truthful $L1$ (approximately) coincides with equilibrium.

To the extent that a truthful $L1$ underbids in a first-price auction with common values, value adjustment makes a truthful $L2$ tend to overbid; but the bidding trade-off may either reinforce or offset this tendency. In general a truthful $L2$ may either underbid or overbid. On balance, it seems more likely to overbid. In KL's example, a truthful $L2$ (approximately) coincides with equilibrium because $L1$ does.

To sum up for first-price auctions, with independent private values level- k types of either kind coincide with equilibrium and cursed equilibrium in the examples that have most often been studied experimentally, where the value distribution is uniform, and so cannot improve upon their explanations of subjects' bidding behavior. But for more general value distributions, as in GHP (2002), a random (or truthful) $L1$ coincides with equilibrium (or underbids) in the low-value treatment and both overbid in the high-value treatment, and a random (or truthful) $L2$ coincides with equilibrium (or underbids) for both treatments. Thus a level- k model is clearly separated from equilibrium and cursed equilibrium, and may be able to explain non-equilibrium bidding.

With common values, a random $L1$ or $L2$ may either underbid or overbid in general. In KL's example when $N = 2$, a random $L1$ and $L2$ both (approximately) coincide with equilibrium and fully-cursed equilibrium; and when $N > 2$, a random $L1$ (approximately) coincides with fully-cursed equilibrium but overbids relative to equilibrium, and a random $L2$ (approximately) coincides with equilibrium and underbids relative to fully-cursed equilibrium. This gives a random level- k model the potential to improve upon χ -cursed-equilibrium explanations of the curse in first-price auctions with common values; but this depends on whether its mixture of $L1$ and $L2$ bidders gives a better account of subjects' heterogeneous bidding behavior than χ -cursed-equilibrium. By contrast, in the only first-price common-value design for which data are

available, KL's, truthful $L1$ and $L2$ bidders coincide with equilibrium, which makes it difficult to reconcile a truthful level- k model with observed bidding behavior.

4d. Cross-treatment implications of cursed-equilibrium versus level- k models

Cursed equilibrium and level- k models also have cross-treatment and aggregate implications, some of which can be tested via between- or within-subjects comparisons using existing experimental data, and some of which are potentially testable in new experiments. We focus on the random level- k model because the truthful level- k model appears to have little potential to explain behavior. We start by discussing the implications type by type, and then translate them into hypotheses about the population type frequencies.

First, a random level- k model predicts more "as-if-equilibrium" play in the aggregate in settings where either more, or more frequent, types' bidding strategies coincide with equilibrium. In KL's example, for instance, a level- k model gives a structural explanation of the heterogeneity of subjects' responses, predicting a stable mixture of as-if-equilibrium ($L2$) and as-if-fully-cursed equilibrium ($L1$) bidders, normally with more of the latter because $L1$ is more prevalent (Costa-Gomes and Crawford (2005)). In KL's example, unlike cursed equilibrium, a random level- k model predicts more as-if-equilibrium play in first- than in second-price common-value auctions because a random $L2$ coincides with equilibrium in the former but not the latter.²⁹

A random level- k model also has different revenue implications than χ -cursed equilibrium, which allows further tests using existing data or data from new experiments. The theoretical prediction for equilibrium and cursed equilibrium is based on MW's equilibrium result that with affiliated signals a second-price auction always yields expected revenue at least as high as a first-price auction (MW, Theorem 15). Assuming χ is the same in first- and second-price auctions, equilibrium's qualitative revenue rankings extend unchanged to χ -cursed equilibrium for any χ : If signals are independent as in AK's example, first- and second-price auctions yield the same expected revenue; and if signals are affiliated as in KL's example, a second-price auction yields expected revenue at least as high (strictly higher in KL's example) as a first-price auction.

By contrast, a random level- k model has revenue implications that depend on type, which for some population type frequencies can weaken or reverse the equilibrium and cursed-equilibrium revenue ranking. For a random $L1$ bidder, a second-price auction yields expected revenue higher than a first-price auction. (This is as expected for KL's example, where a random

²⁹Cursed equilibrium could replicate such a prediction by estimating subject-specific χ parameters with $\chi = 0$ or 1 for individual subjects; but this would give it many more free parameters than a level- k model.

L1 coincides with fully-cursed equilibrium, which implies the same revenue ranking as equilibrium.) But for a random *L2* bidder, this ranking is reversed when $N > 2$: In a second-price auction, *L2* bidders then underbid relative to equilibrium; but in a first-price auction they coincide with equilibrium, and so yield higher expected revenue. Thus, if the population has a roughly equal mix of *L1* and *L2* bidders, which is empirically plausible, a random level- k model will tend to predict a weaker revenue ranking than equilibrium or even fully-cursed equilibrium; and if there are more *L2* than *L1* bidders, which is not implausible for experienced bidders, a random level- k model will reverse the equilibrium and cursed-equilibrium ranking.

Finally, a random level- k model, like cursed equilibrium, can explain why the curse gets worse and expected revenue increases with higher N (ER (2005, Proposition 6)). In a second-price auction a random *L1* coincides with a fully-cursed equilibrium bidder, so to the extent that *L1* predominates a level- k model will have similar implications about the effects of increasing N . ER (2005, Propositions 7 and 8) show that when $r(x)$ is a symmetric random variable then bidders will have negative expected payoffs in a second-price auction as long as $N > 3$, and that the effect of N on expected revenue increases with χ . Because in a second-price auction a random *L1*'s bids coincide with fully-cursed equilibrium bids, these results hold for *L1* decisions too.

5. Comparing the Models in Auction Experiments

With the exception of equilibrium, all the models compared here depend on behavioral parameters: population type frequencies for level- k models, one or more cursedness parameters for cursed equilibrium, and one or more logit precisions for QRE. This section makes our analysis more concrete by using data from auction experiments previously gathered by others to estimate these parameters econometrically, and then using the results to compare the models' abilities to account for observed behavior in the experiments. Our goal in using econometric estimates is to constrain our discretion in calibrating the models and to obtain likelihoods that provide an objective criterion for comparing them; not to take a definitive position on the parameters.

Table 2 summarizes the data used in our analysis. The data were chosen with two goals in mind. First, because learning can lead even unsophisticated subjects to equilibrium, strategic thinking appears most clearly before subjects have seen other subjects' responses; we therefore use data from inexperienced subjects. Given this, we maximize comparability with ER's analysis of common-value auction data from KL (1986) for first-price auctions and AK (1997) for second-price auctions. However, KL had only experienced subjects (defined as those who had

participated in at least one prior auction series); and while AK had some inexperienced subjects, ER's analysis focused mainly on their experienced subjects. For common-value second-price auctions we use AK's data for inexperienced subjects and the unpublished data for inexperienced subjects in a second-price version of KL's design mentioned in the Appendix to Kagel, Levin, Battalio, and Meyer (1989) as reprinted in KL (2002, Chapter 2). For common-value first-price auctions we use Garvin and Kagel's (1994) data for inexperienced subjects in KL's design. Further, instead of pooling the data from all periods and usually all subjects as ER did, we focus on individual subjects' initial responses, interpreted as the first 5 periods (in which a subject typically had 5 different realizations of his value signal) to compensate for the small sample size.

Finally, because cursed equilibrium coincides with equilibrium in independent-private-value auctions, they are particularly important in assessing the predictive value of the level- k model's more general view of strategic behavior. But with independent private values, all level- k types (truthful or random) coincide with equilibrium in second-price auctions; and with the i.i.d. uniform values used in most designs, in first-price auctions as well (Section 4). We therefore use GHP's (2002) data from independent-private-value first-price auctions with discrete, slightly non-uniform values, in which level- k types are separated from equilibrium.³⁰

Our econometric specification follows the models with mixtures of decision rules or *types* of Stahl and Wilson (1994, 1995), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004), Costa-Gomes and Crawford (2004), and Crawford and Iriberry (2005). Index Table 2's experimental treatments (or "games") $g = 1, 2, 3, 4$; and game g 's subjects $i = 1, 2, \dots, N_g$.

For our level- k plus equilibrium model we allow random $L0$ and both random and truthful $L1$ and $L2$ as well as *Equilibrium*, treating each as a separate type with its own beliefs as specified in Section 2 or 3, indexing them $k = 1, 2, \dots, K$.³¹ For our cursed-equilibrium model we allow levels of cursedness χ that vary across subjects (as in part of ER's analysis of AK's data); but we constrain χ to multiples of 0.1 in the interval $[0, 1]$ and in some cases to an estimated subset of

³⁰Other common-value settings whose data would enrich our analysis include LP's (1991) and HS's (2000), but despite their authors' generous efforts, those data are not yet available. We define payoffs as payments for performance, exclusive of show-up fees, etc.; and express all payoffs in 1989 dollars. Following AK and GHP, we edited a small number of "crazy" bids (6 in AK (1997), 11 in KL (1986) first-price, 3 in KL (1986) second-price, and 12 in GHP), replacing bids above or below the highest or lowest rationalizable bid with the highest or lowest rationalizable bid, respectively.

³¹We omit truthful $L0$ in the econometric analysis because consistently truthful bidding is very rare in the data for the first-price treatments (6/255 observations in KL and 6/400 in GHP, with no individual subject making more than two truthful bids); and because there is no way to assign beliefs that makes truthful bidding optimal in first-price auctions, where it is dominated, which makes it difficult to specify logit errors like those we use for the other types.

possible χ 's in the interval $[0, 1]$ (with no constraint to multiples of 0.1), whose size is chosen to make the model's number of parameters more comparable to that of our level- k plus equilibrium model.³² We then treat each possible level of χ as a separate type with its own beliefs as specified in Section 2, indexed $k = 1, 2, \dots, K$. Either way, π_k denotes the proportion of type k in the population, with $\sum_k \pi_k = 1$. Thus each model includes an *Equilibrium* type ($\chi = 0$ for cursed equilibrium), which allows a fair comparison of how well they explain subjects' deviations from equilibrium. Our formal discussion covers both models.

Type k implies a bidding strategy in game g (whether it is first- or second-price) denoted $c_k^g(x)$; and we write c_{it}^g for subject i 's observed bid in game g at time t . We assume that a subject of type k normally follows k 's prescribed bidding strategy $c_k^g(x)$, but subject to logistic errors, which are independent across the five periods in which he plays. Let the density $f_k^g(z)$ represent the beliefs about the highest bid among the others, Z_1 , implicit in type k in game g . Letting $S_g(c, x, z)$ denote a player's expected payoff in game g for bid c , signal x , and highest bid among the others z , a subject's expected payoff in game g for type k 's beliefs with signal x can be written:

$$(34) \quad S_k^g(c | x) \equiv \int_c^{c_{k-1}^{-1}(c_k)} S_g(c, x, z) f_k^g(z) dz .$$

Given a subject of type k with signal x and precision λ in game g , the probability of observing bid c within the range of possible bids $[\underline{c}, \bar{c}]$ is then:

$$(35) \quad \Pr(c | k, x, g, \lambda) = \frac{\exp(\lambda S_k^g(c | x))}{\int_{\underline{c}}^{\bar{c}} \exp(\lambda S_k^g(e | x)) de} .$$

Thus, the costlier an error is ex ante, given type- k beliefs, the lower the probability of making it, with the cost-sensitivity tuned by the precision parameter λ and bids approaching uniform randomness as $\lambda \rightarrow 0$ or the error-free bid specified by $c_k^g(x)$ as $\lambda \rightarrow \infty$.

With errors independent, conditional on type, the likelihood of observing the 5-observation sample $c_i^g = (c_{i1}^g, c_{i2}^g, c_{i3}^g, c_{i4}^g, c_{i5}^g)$ for subject i of type k with signal x and precision λ in game g is:

³²ER (2005, Table II), by contrast, allow χ to take any real value and report many estimates for AK's inexperienced subjects outside $[0, 1]$, despite χ 's interpretation as the probability a player assigns to others playing their average distribution of actions irrespective of type rather than their type-contingent strategies.

$$(36) \quad L_k(c_i^g | k, x, g, \lambda) = \prod_{t=1}^5 \Pr(c_{it}^g | k, x, g, \lambda).$$

The likelihood of observing c_i^g unconditional on type is:

$$(37) \quad \sum_{k=1}^K \pi_k L_k(c_i^g | k, x, g, \lambda) = \sum_{k=1}^K \pi_k \prod_{t=1}^5 \Pr(c_{it}^g | k, x, g, \lambda).$$

Given (35), the likelihood in (37) treats a bid as stronger evidence for a type the closer it is to the type's optimal bid or the better the deviations are explained given the type's beliefs, because the payoff function is quasiconcave and the logit term increases with payoff. In most cases the types are well separated and the first factor is the dominant one, but in GHP random *LI* and *Equilibrium* bids are separated only for $v = 12$ in the high-value treatment, and subjects are separated mainly by the differences in deviation costs implied by their beliefs.

To allow for subject heterogeneity, we compare three specifications of how the precision λ varies by subject and/or type. Let the matrix $\Lambda \equiv [\lambda_{ik}]$ denote the precisions indexed by subject i and type k . Subject-specific error precisions, the most flexible specification, places no restrictions on how λ varies with i and k . Type-specific error precisions restricts λ_{ik} to be independent of i for any given k . Constant error precisions restricts λ_{ik} to be independent of i and k . In each case, the precision is the same for all of a given subject's bids, so that (35)-(37) are well-defined as written.

Letting $c^g = (c_1^g, c_2^g, \dots, c_{N_g}^g)$, from (37) we can now write the likelihood and log-likelihood functions for game g :

$$(38) \quad L(\pi, \Lambda | c^g) = \prod_{i=1}^{N_g} \sum_{k=1}^K \pi_k L_k(c_i^g | k, x, g, \lambda_{ik}) \text{ and } LL(\pi, \Lambda | c^g) = \sum_{i=1}^{N_g} \log \left(\sum_{k=1}^K \pi_k L_k(c_i^g | k, x, g, \lambda_{ik}) \right)$$

As explained above, (34)-(38) define our cursed-equilibrium model as well as our level- k plus equilibrium model, interpreting each type k as a different level of the cursedness parameter χ . The levels of cursedness are sometimes fixed and sometimes estimated, as explained below. To assure comparability with our level- k model, cursed types are assumed to make logit errors with analogous specifications of how their precisions vary by subject and/or type.

Tables 3a-d summarize treatment-by-treatment parameter estimates and likelihoods for the level- k plus equilibrium and cursed-equilibrium models and, in Table 3d, QRE for GHP.³³ For the

³³We have not tried to estimate the models for all treatments together due to widely differing subject pools and conditions. In GHP we define random *L0* with equal probabilities for the possible values in each treatment.

level- k plus equilibrium model we include all the types discussed in Section 3, with those that are not separated in a given treatment included in the table and indicated by a tilde (\sim).

Individual subjects' precisions are highly heterogeneous. Likelihood-ratio tests for the level- k plus equilibrium model, for which our three alternative specifications are nested, strongly reject constant or type-specific error precisions in all four treatments, with p -values of 0.001 or lower. The Akaike and Bayesian Information Criteria both favor a specification with subject-specific precisions, except for GHP where they favor a model with constant precisions. For our cursed-equilibrium model the alternative specifications are not nested, but the Information Criteria both favor a specification with subject-specific precisions. For simplicity, we focus on the results with subject-specific precisions for both models. The estimation in (38) is then equivalent to estimating subject-by-subject, and the possible dependence of precision on type is redundant.

The estimated population type frequencies for the level- k plus equilibrium model are generally behaviorally plausible. As in other settings, the estimated frequency of random $L0$ is 0 in three treatments and negligible in the fourth, suggesting that the non-strategic anchoring type exists mainly in the minds of random $L1$ and random $L2$; and that subjects' behavior, while not usually conforming to equilibrium, is nonetheless strategic. The estimates for KL first-price, AK second-price, and GHP, particularly with subject-specific precisions, are very close to each other and (identifying *Equilibrium* with random $L2$ in KL first-price) close to previous estimates from other settings (Stahl and Wilson (1995), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004), Costa-Gomes and Crawford (2004), and Crawford and Iriberri (2005)).

The estimates for KL second-price are quite different, and to us less plausible. In that treatment *Equilibrium* shades its bid below the value suggested by its signal to adjust for the curse, random $L1$ bids the value suggested by its signal, random $L2$ shades more than equilibrium, and random $L3$ or truthful $L2$ (which are not separated) bid above the value suggested by their signal. There are two main patterns in the data: First, some subjects shade their bids, but less than in equilibrium; in our model this is best captured by *Equilibrium* or random $L1$. Second, other subjects bid above the value suggested by their signal, which is best captured by random $L3$ or truthful $L2$. We suspect that these subjects bid so high not because they believe others are shading their bids more than they would in equilibrium (as random $L3$ and truthful $L2$ believe); but rather because they don't fully process the subtle implications of the second-price auction for their optimal bidding strategy: Because they know they will not have to pay their own bid, they

underestimate its indirect cost, which may be less salient to them than what they will pay. If the random $L3$ /truthful $L2$ type is omitted, these subjects are best described by random $L1$.

Turning to cursed equilibrium, for the three common-value treatments we begin with the model in which χ is constrained to multiples of 0.1 in the interval $[0, 1]$, and again focus on the estimates with subject-specific precisions. Although this model has many more parameters than the corresponding level- k plus equilibrium model, it is useful for diagnostic purposes.³⁴

Despite our different specification and our use of data from inexperienced subjects, our cursed-equilibrium estimates for KL's and AK's designs are generally consistent with ER's estimates for KL's and AK's subjects, particularly AK's inexperienced subjects.³⁵ And our cursed-equilibrium estimates are very close to our estimates for the level- k plus equilibrium model: In the cursed-equilibrium estimates for all three common-value treatments, there are two spikes in the distribution of individual estimates, one at $\chi = 0$ (*Equilibrium*) and one at $\chi = 1$ (fully-cursed equilibrium or random $L1$), with comparatively little probability mass in between (with minor exceptions at $\chi = 0.2$ in KL second-price and $\chi = 0.7$ in AK second-price). Intermediate cursed types add little explanatory power over level- k types.

Even when χ is allowed to take any value that is a multiple of 0.1 in the interval $[0, 1]$, with subject-specific precisions the level- k plus equilibrium model has a likelihood advantage over the cursed-equilibrium model for KL and AK second-price; but in this case the cursed-equilibrium model has a likelihood advantage for KL first-price. In each case, the cursed-equilibrium model has many more parameters, so we also do the comparisons with a cursed-equilibrium model whose number of parameters equals that of the level- k plus equilibrium model. If we allow two cursed-equilibrium types in KL first-price, say $\chi = 0$ and $\chi = 1$ as the estimates when χ is constrained to multiples of 0.1 in $[0, 1]$ suggest, then with subject-specific precisions the level- k model (-1660.52) has a small likelihood advantage over the cursed-equilibrium model (-1663.85). Since in this case both models' types have the same optimal bidding strategies, the level- k model's

³⁴The number of estimated parameters for the level- k plus equilibrium model is $N + K - 1$, where K is the number of separated types in the treatment, and so varies from 52 in KL first-price to 31 in KL second-price, 26 in AK second-price, and 84 in GHP. By contrast, the cursed-equilibrium model with subject-specific precisions has $10 + N$ parameters in the common-value treatments, or 61 in KL first-price, 38 in KL second-price, and 33 in AK second-price. (The number of observations is $5N$ in each treatment.)

³⁵In KL's and AK's designs, cursed equilibrium bids are linear in both the bidder's private signal x and the cursedness parameter χ . Pooling the data across time periods, ER regressed subjects' bids on those variables, finding that when constrained to be equal for all subjects, χ is closer to 1 (fully-cursed equilibrium) for inexperienced subjects and to 0 (equilibrium) for experienced subjects; and that for AK's data, when χ was allowed to vary across subjects, it varied much more for inexperienced than experienced subjects, and was significantly different from 0 for both.

advantage stems from random *LI*'s slight likelihood advantage over fully-cursed equilibrium, given random *LI*'s non-equilibrium beliefs, in describing deviations from the optimal bidding strategy. However, in KL first-price the constraint that $\chi = 0$ or $\chi = 1$ is strongly rejected, and allowing more intermediate levels of χ allows the cursed-equilibrium model to fit some subjects better than the level- k types, with a better overall fit as well.³⁶

Turning to GHP's independent-private-value design, we replace cursed equilibrium, which is not separated from equilibrium, by a mixture model that allows two different QRE "types" with different, estimated precisions. Here we follow GHP's preferred explanation of their subjects' non-equilibrium bids, but we again restrict the number of types to make the number of parameters comparable to that of the level- k plus equilibrium model. Again focusing on the estimates with subject-specific precisions, the level- k model has a likelihood advantage over a QRE model.³⁷ Random *LI* dominates the level- k type estimates, but there is a significant number of *Equilibrium* subjects and a smaller but significant number of truthful *LI* subjects, the latter mostly from the high-value treatment, where truthful *LI* predicts the overbidding that occurs for some values.

6. Conclusion

This paper has proposed a new approach to explaining the winner's curse in common-value auctions and overbidding in some independent-private-value auctions, based on a structural non-equilibrium "level- k " model of initial responses that describes behavior in a variety of experiments with complete-information games. We consider alternative ways to generalize complete-information level- k models to this leading class of incomplete-information games, and derive their implications in first- and second-price auctions with general information structures, comparing them to equilibrium and Eyster and Rabin's (2005) notion of "cursed equilibrium."

Our analysis shows that many of the insights of equilibrium auction theory, properly interpreted, are robust to empirically plausible failures of the equilibrium assumption. It yields a tractable non-equilibrium characterization of the value adjustment for the information revealed by winning (the "winner's curse") that influences equilibrium bidding strategies in first- or second-price common-value auctions; and of the bidding trade-off between the cost of higher bids and

³⁶In the specifications with type-specific or constant precisions, we also allowed the same number of cursed types as level- k types. But in these cases, for KL first-price and AK second-price, the cursed-equilibrium model estimated fewer than the number of types we allowed it. The same thing happened with the QRE types estimated for GHP.

³⁷Although the Akaike and Bayesian Information Criteria favor the model with constant precision for GHP, its type frequency estimates do not differ significantly from those with subject-specific precisions. Random *L2* is separated from *Equilibrium* only in GHP's high-value treatment, and then only weakly.

their higher probability of winning that influences equilibrium bidding strategies in first-price auctions with common or independent private values. These characterizations guide the choice of a model that can track the patterns of variation in subjects' behavior across different treatments with an empirically plausible population distribution of level- k types. By allowing us to examine auction behavior through the lens of a more general model of strategic behavior, they also allow us to link a large body of data from auction experiments to data from experiments in other settings that were specifically designed to explore strategic thinking.

In our econometric analysis, we find that a specification based on a random uniform anchoring (level-0) type like the one used in many analyses of complete-information games allows a unified explanation of the winner's curse in common-value auctions and overbidding in those independent-private-value auctions without the uniform value distributions used in most experiments; but that a specification based on a truthful anchoring type, despite its plausibility in auctions, does poorly either because it is not separated from equilibrium, or it requires an empirically implausible type distribution. A level- k model fits subjects' initial responses to a wide range of auction experiments better than the leading alternatives of cursed equilibrium or (in one case) quantal-response equilibrium.

References

- Avery, Christopher and Kagel, John H. (1997) "Second-Price Auctions with Asymmetric Payoffs: An Experimental Investigation." *Journal of Economics and Management Strategy*, 6(3), 573–603.
- Ball, Sheryl; Bazerman, Max H.; and Carroll, John (1991). "An Evaluation of Learning in the Bilateral Winner's Curse," *Organizational Behavior and Human Decision Processes* 48, 1-22.
- Battigalli, Pierpaolo, and Siniscalchi, Marciano (2003). "Rationalizable Bidding in First-Price Auctions," *Games and Economic Behavior* 45, 38-72.
- Bazerman, Max H. and Samuelson, William F. (1983). "I Won the Auction But I Don't Want the Prize," *Journal of Conflict Resolution* 27(4), 618-634.
- Cai, Hongbin; and Wang, Joseph Tao-Yi (2005). "Overcommunication in Strategic Information Transmission Games," *Games and Economic Behavior*, in press.
- Camerer, Colin; Ho, Teck-Hua; and Chong, Juin-Kuan (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119, 861-898.

- Capen, E. C.; Clapp, R. V.; and Campbell, W. M. (1971). "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 23(6), 641–53.
- Charness, Gary; and Levin, Dan (2005): "The Origin of the Winner's Curse: A Laboratory Study," manuscript.
- Chen, Kay-Yut; and Plott, Charles R. (1998). "Nonlinear Behavior in Sealed Bid First Price Auctions," *Games and Economic Behavior*, 25, 34-78.
- Costa-Gomes, Miguel; Crawford, Vincent; and Broseta, Bruno (2001): "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 69, 1193-1235.
- Costa-Gomes, Miguel; and Crawford, Vincent P. (2004): "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," UCSD Discussion Paper 2004-11, April 2004, revised October 2004, <http://dss.ucsd.edu/~vcrawfor/#Guess>.
- Cox, J. C.; Smith, V. L.; and Walker, J. M. (1983), *Economics Letters* , 12, 207-212.
- Cox, J. C.; Smith, V. L.; and Walker, J. M. (1988), "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, 1, 61–99.
- Cox, J. C.; Smith, V. L.; and Walker, J. M. (1992), "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review* 82, 1392–1412.
- Crawford, Vincent P. (1998), "A Survey of Experiments on Communication via Cheap Talk," *Journal of Economic Theory* 78, 286-298
- Crawford, Vincent P. (2003): "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review*, 93, 133-149.
- Crawford, Vincent P.; and Iriberry, Nagore (2005): "Fatal Attraction: Focality, Naivete, and Sophistication in Experimental 'Hide-and-Seek' Games," UCSD Discussion Paper 2004-12R, revised January 2005, <http://dss.ucsd.edu/~vcrawfor/#Hide>.
- Dyer, Douglas; Kagel, John H.; and Levin, Dan (1989), "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis," *The Economic Journal*, 99 (394), 108-115.
- Eyster, Erik; and Rabin, Matthew (2002). "Cursed Equilibrium," IBER, UC Berkeley Working Paper E02-320, August 2002.
- Eyster, Erik; and Rabin, Matthew (2005). "Cursed Equilibrium," *Econometrica* 73(5), 1623-1672.
- Feddersen, Timothy; and Pesendorfer, Wolfgang (1996). "The Swing Voter's Curse," *American Economic Review* 86(3), 408-424.

- Feddersen, Timothy; and Pesendorfer, Wolfgang (1997). "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica* 65(5), 1029-1058.
- Feddersen, Timothy; and Pesendorfer, Wolfgang (1998). "Convicting the Innocent: the Inferiority of Unanimous Jury Verdicts under Strategic Voting," *American Political Science Review*, 90(1), 23-35.
- Friedman, Daniel (1998), "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly," *American Economic Review* 88:4, 933-946.
- Garvin, Susan; and Kagel, John H. (1994), "Learning in Common Value Auctions: Some Initial Observations," *Journal of Economic Behavior & Organization*, 25(3), 351-372.
- Goeree, Jacob K.; Holt, Charles A.; and Palfrey, Thomas R. (2002), "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104 (1), 247-272.
- Goeree, Jacob K.; and Offerman, Theo (2002a). "Efficiency in Auctions with Private and Common Values: An Experimental Study," *American Economic Review* 92(3), 625-643.
- Goeree, Jacob K.; and Offerman, Theo (2002b), "Winner's Curse Without Overbidding," manuscript, University of Virginia.
- Harrison, Glenn W. (1989). "Theory and Misbehavior in First-Price Auctions," *American Economic Review*, 79: 749-62.
- Hendricks, Kenneth; Porter, Robert H.; and Boudreau, Bryan (1987). "Information, Returns, and Bidding Behavior in OCS Auctions: 1954-1969," *Journal of Industrial Economics*, 35, 517-42.
- Hendricks, Kenneth; and Porter, Robert H. (1988), "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 78, 865-883.
- Ho, Teck-Hua; Camerer, Colin; and Weigelt, Keith (1998): "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'," *American Economic Review*, 88, 947-969.
- Holt, Charles A. (1980), "Competitive Bidding for Contracts under Alternative Auction Procedures," *Journal of Political Economy*, 88: 433-45.
- Holt, Charles A.; and Sherman, Roger (1994). "The Loser's Curse," *American Economic Review*, 84(3), 642-652.
- Holt, Charles A.; and Sherman, Roger (2000). "Risk Aversion and the Winner's Curse." Working paper, University of Virginia.

- Kagel, John H. (1995). "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, John H. Kagel and Alvin E. Roth, eds, Princeton NJ: Princeton University Press, 1995.
- Kagel, John H.; and Levin, Dan (1986). "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76(5), 894-920.
- Kagel, John H.; Levin, Dan; Battalio, Ray C.; and Meyer, D. J. (1989). "First-Price Common Value Auctions: Bidder Behavior and the 'Winner's Curse'," *Economic Inquiry*, 27, 241-258.
- Kagel, John H.; and Levin, Dan (2002). *Common Values Auctions and the Winner's Curse*. Princeton, NJ: Princeton University Press.
- Kagel, John H.; and Jean-Francois Richard (2001). "Super-Experienced Bidders in First-Price Common Value Auctions: Rules of Thumb, Nash Equilibrium Bidding, and the Winner's Curse," *Review of Economics and Statistics*, 83, 408-419.
- Kagel, John H.; and Roth, Alvin (1992), "Theory and Misbehavior in First-Price Auctions: Comment," *American Economic Review*, 82: 1379-1391.
- Klemperer, Paul (2000), "Auction Theory: A Guide to the Literature," chapter 1 of *The Economic Theory of Auctions*, Paul Klemperer, ed., Edward Elgar, Vol. 1, 2004.
- Laffont, Jean-Jacques (1997), "Game Theory and Empirical Economics: The Case of Auction Data," *European Economic Review*, 41(1), 1-35.
- Lind, Barry; and Plott, Charles R. (1991). "The Winner's Curse: Experiments with Buyers and with Sellers." *American Economic Review*, 81(1), 335-46.
- Maskin, Eric; and Riley, John (2000), "Asymmetric Auctions," *Review of Economic Studies* 67, 413-438.
- McAfee, R. Preston; and McMillan, John (1987). "Auctions and Bidding," *Journal of Economic Literature*, 25 (2), 699-738.
- McKelvey, Richard; and Palfrey, Thomas (1995). "Quantal Response Equilibria for Normal-Form Games," *Games and Economic Behavior*, 10, 6-38.
- Milgrom, Paul; and Weber, Robert (1982), "A Theory of Auctions and Competitive Bidding", *Econometrica*, 50, 1089-1122.

- Milgrom, Paul (1985), "The Economics of Competitive Bidding: A Selective Survey," ch. 9, pp. 261-289 in *Social Goals and Social Organization: A Volume in Honor of Elisha Pazner*, eds. L. Hurwicz, D. Schmeidler and H. Sonnenschein, London: Cambridge University Press, 1985.
- Milgrom, Paul (1987), "Auction Theory," pp. 1-32 in *Advances in Economic Theory: Fifth World Congress*, ed. Truman Bewley, London: Cambridge University Press, 1987.
- Milgrom, Paul (1989). "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, 3, 3-22.
- Nagel, Rosemarie (1995): "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313-1326.
- Palacios-Huerta, Ignacio (2003). "Learning to Open Monty Hall's Doors," *Experimental Economics*, 6, 235–251.
- Palfrey, Thomas (1985). "Buyer Behavior and the Welfare Effects of Bundling by a Multiproduct Monopolist: A Laboratory Test," in *Research in Experimental Economics 3*, ed. Vernon Smith, JAI Press, pp. 73-104.
- Porter, Robert H. (1995), "The Role of Information in U.S. Offshore Oil and Gas Lease Auctions," *Econometrica*, 63: 1-27.
- Samuelson, William F.; and Bazerman, Max H. (1985). "The Winner's Curse in Bilateral Negotiations." In *Research in Experimental Economics*, vol. 3, Vernon L. Smith, ed., Greenwich, CT: JAI Press, 1985.
- Selten, Reinhard (1998). "Features of Experimentally Observed Bounded Rationality," *European Economic Review*, 42, 413-436.
- Sonsino, Doron; Erev, Ido; and Gilat, Sharon (2002). "On the Likelihood of Repeated Zero-Sum Betting by Adaptive (Human) Agents," manuscript, Technion, Israel Institute of Technology.
- Sovik, Ylva (2000). "Impossible Bets: An Experimental Study," manuscript, University of Oslo, Department of Economics.
- Stahl, Dale; and Wilson, Paul (1994): "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization*, 25, 309-327.
- Stahl, Dale; and Wilson, Paul (1995): "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 218-254.

- Thaler, Richard H. "The Winner's Curse," *Journal of Economic Perspectives* 2(1), 191-202.
- Tor, Avishalom; and Bazerman, Max H. (2003). "Focusing Failures in Competitive Environments: Explaining Decision Errors in the Monty Hall Game, the Acquiring a Company Problem, and Multiparty Ultimatums," *Journal of Behavioral Decision Making*, 16, 353-374.
- Vickrey, William (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8-37.
- Weizsäcker, Georg (2003): "Ignoring the Rationality of Others: Evidence from Experimental Normal-Form Games," *Games and Economic Behavior*, 44, 145-171.
- Wilson, Robert (1969). "Competitive Bidding with Disparate Information," *Management Science*, 15(7), 446-448. Reprinted in *The Economics of Information*, eds. Steven A. Lippman and David K. Levine, Edward Elgar, London, 1994; and in *The Economic Theory of Auctions*, ed. Paul Klemperer, Edward Elgar, London, 1999.
- Wilson, Robert (1992) "Strategic Analysis of Auctions," ch 8, pp. 227-279 in *Handbook of Game Theory*, Volume 1, eds. Robert Aumann and Sergiu Hart, Amsterdam: North-Holland/Elsevier Science Publishers.

Table 1. Types' Bidding Strategies^a

Auction/Type	Equilibrium	χ-cursed Equilibrium	Random $L1$	Random $L2$	Truthful $L1$	Truthful $L2$
2 nd -price i.p.v.	x	x	$b_1^r(x) = x$	$b_2^r(x) = x$	$b_1^t(x)$ from (30) with $v(x, \cdot) \equiv x$	$b_2^t(x)$ from (32) with $v(x, \cdot) \equiv x$
2 nd -price c.v.	$b_*(x) = v(x, x)$	$b_\chi(x) = (1-\chi)v(x, x) + \chi r(x)$	$b_1^r(x) = r(x)$	$b_2^r(x)$ from (28): $b = v(x, b_1^{r^{-1}}(b))$	$b_1^t(x)$ from (30): $b = v(x, r^{-1}(b))$	$b_2^t(x)$ from (32): $b = v(x, b_1^{t^{-1}}(b))$
2 nd -price c.v.: KL	$x - \frac{a}{2} + \frac{a}{N}$	$x - (1-\chi)a \frac{N-2}{2N}$	x	$x - \frac{a}{2} \left(\frac{N-2}{N-1} \right)$	$x - \frac{a}{2} \left(\frac{N-2}{N-1} \right)$	We have not derived a solution, but bounds are derived in the Appendix
2 nd -price c.v.: AK	$2x$	$\chi \left(x + \frac{5}{2} \right) + (1-\chi)2x$	$x + \frac{5}{2}$	3.5 if $x \leq 2.5$; 6.5 if $x > 2.5$	3.5 if $x \leq 2.5$; 6.5 if $x > 2.5$	We have not derived a solution, but bounds are derived in the Appendix
1 st -price i.p.v.	$a_*(x)$ from (4)	$a_*(x)$ from (4)	$a_1^r(x)$ from (15)	$a_2^r(x)$ from (27) with $v(x, \cdot) \equiv x$	$a_1^t(x)$ from (29) with $v(x, \cdot) \equiv x$	$a_2^t(x)$ from (31) with $v(x) = x$
1 st -price c.v.	$a_*(x)$ from (3)	$a_\chi(x)$ from (11)	$a_1^r(x)$ from (15)	$a_2^r(x)$ from (27)	$a_1^t(x)$ from (29)	$a_2^t(x)$ from (31)
1 st -price c.v.: KL	$x - \frac{a}{2}$ $+ \frac{a}{N} \exp\left(-\frac{N(x-x)}{a}\right)$	$\left[\chi + (1-\chi) \left(x - \frac{a}{2} + \frac{a}{N} \right) - \frac{a}{N} \right]$ $+ \frac{a}{N} \exp\left(-\frac{N(x-x)}{a}\right)$	$x - \frac{a}{N}$	$x - \frac{a}{2}$	$x - \frac{a}{2}$	$x - \frac{a}{2}$

^aIf there is no general closed-form expression, Table 1 refers to the equation in the text that determines the bidding strategy.

Table 2. Data Sources and Experimental Designs					
<i>g</i> (Experimental treatment)	Auction Type	$u(S, X)$	Signals	n (sample size)	Treatment variables
1. KL (1986)	First-Price Common Value	$u(S, X) = S$	$X S \sim U[s - \frac{a}{2}, s + \frac{a}{2}]$	51	a (precision), N (number of bidders), limits of s
2. KL (1986)	Second-Price Common Value	$u(S, X) = S$	$X S \sim U[s - \frac{a}{2}, s + \frac{a}{2}]$	28	a (precision)
3. AK (1997)	Second-Price Common Value	$u(S, X) = X_1 + X_2$	$X \sim U[\underline{x}, \bar{x}] = [1, 4]$	23	No variation, $N = 2$
4. GHP (2002)	First-Price Independent Private Value	$u(S, X) = X$	$X \sim U[0, 2, 4, 6, 8, 11]$ $X \sim U[0, 3, 5, 7, 9, 12]$	40	No variation, $N = 2$

Table 3a. Estimation Approaches and Estimates for Kagel and Levin First-Price

Model	Level- k plus equilibrium			Mixture of cursed types							
	Individual specific precision (λ_i)	Type specific precision (λ_k)	Same precision (λ)	Individual specific precision (λ_i) and fixed cursedness types ($\chi = (1, 0.9, \dots, 0)$)				Type specific precision (λ_k)	Same precision (λ)	Type specific precision (λ_k)	Same precision (λ)
Specification				Types	χ		Reduced types	χ_k		χ_k	
<i>Random L0</i> ($\lambda = 0$)	0.04	--	--	<i>RL0</i> ($\lambda = 0$)		0.06	0.04	--		--	
<i>Random L1</i>	0.67	0.89	0.51	Type 1	1	0.47	0.67	0.84	0.89	1	0.5
λ_{RL1}		2.24	1.31	Type 2	0.9	0.02	--	λ_1	0.59	λ_2	0.68
<i>Random L2</i>	<i>~Equilibrium</i> <i>m</i>	<i>~Equilibrium</i>	<i>~Equilibrium</i>	Type 3	0.8	0.08	--	0	0.11	0	0.5
λ_{RL2}				Type 4	0.7	0.06	--	λ_2	14.84	λ_1	0.68
<i>Truthful L1</i>	<i>~Equilibrium</i> <i>m</i>	<i>~Equilibrium</i>	<i>~Equilibrium</i>	Type 5	0.6	0	--				
λ_{TL1}				Type 6	0.5	0	--				
<i>Truthful L2</i>	<i>~Equilibrium</i> <i>m</i>	<i>~Equilibrium</i>	<i>~Equilibrium</i>	Type 7	0.4	0.04	--				
λ_{TL2}				Type 8	0.3	0.04	--				
<i>Equilibrium</i>	0.29	0.11	0.49	Type 9	0.2	0.04	--				
λ_{EQ}		14.75	1.31	Type 10	0.1	0	--				
				Type 11	0	0.19	0.29				
Log-likelihood	-1660.52	-1754.28	-1755.28			-1640.5	-1663.85		-1745.7		-1762.24

Table 3b. Estimation Approaches and Estimates for Kagel and Levin Second-Price

Model	Level-k plus equilibrium			Mixture of cursed types						
Specification	Individual specific precision (λ_i)	Type specific precision (λ_k)	Same precision (λ)	Individual specific precision (λ_i) and fixed cursedness types ($\chi = (1, 0.9, \dots, 0)$)			Type specific precision (λ_k)		Same precision (λ)	
				Types	χ		χ_k		χ_k	
<i>Random L0</i> ($\lambda = 0$)	0	--	--	<i>RL0</i> ($\lambda = 0$)	--	0.18			--	--
<i>Random L1</i>	0.21	0.10	0.62	Type 1	1	0.18	0.86	0.27	0.79	0.43
λ_{RL1}		95.84	8.91	Type 2	0.9	0.04	λ_1	8.89	λ_3	2.95
<i>Random L2</i>	0.14	0.27	0.11	Type 3	0.8	0.03	0.18	0.3	0.33	0.15
λ_{RL2}		2.5	8.91	Type 4	0.7	0	λ_2	5.35	λ_2	2.95
<i>Truthful L1</i>	<i>~RandomL2</i>	<i>~RandomL2</i>	<i>~RandomL2</i>	Type 5	0.6	0	--	0.43	0	0.42
λ_{TL1}				Type 6	0.5	0.03	λ_3	0	λ_1	2.95
<i>Truthful L2</i>	0.29	0.33	0.27	Type 7	0.4	0.03				
λ_{TL2}		6.1	8.91	Type 8	0.3	0				
<i>Equilibrium</i>	0.36	0.30	0	Type 9	0.2	0.11				
λ_{EQ}		49.76	--	Type 10	0.1	0.04				
				Type 11	0	0.36				
Log-likelihood	-921.73	-967.8	-973.81			-954.5		-987.48		-995.59

Table 3c. Estimation Approaches and Estimates for Avery and Kagel Second-Price

Model	Level- k plus equilibrium			Mixture of cursed types						
	Individual specific precision (λ_i)	Type specific precision (λ_k)	Same precision (λ)	Individual specific precision (λ_i) and fixed cursedness types ($\chi = (1,0.9,\dots,0)$)			Type specific precision (λ_k)		Same precision (λ)	
Specification				Types	χ		χ_k		χ_k	
<i>Random L0</i> ($\lambda = 0$)	0	--	--	<i>RL0</i> ($\lambda = 0$)	--	0.13			--	--
<i>Random L1</i>	0.65	0.58	0.94	Type 1	1	0.44	1	0.37	0.8	1
λ_{RL1}		12.42	4.3	Type 2	0.9	0	λ_1	9.67	λ_1	2.77
<i>Random L2</i>	0.09	0	0.05	Type 3	0.8	0	0.73	0.08		
λ_{RL2}		--	4.3	Type 4	0.7	0.13	λ_1	161.45		
<i>Truthful L1</i>	<i>~RandomL2</i>	<i>~RandomL2</i>	<i>~RandomL2</i>	Type 5	0.6	0.04	0.63	0.55		
λ_{TL1}				Type 6	0.5	0.09	λ_1	1.33		
<i>Truthful L2</i>	0.22	0.42	0	Type 7	0.4	0.04				
λ_{TL2}		1.07	4.3	Type 8	0.3	0				
<i>Equilibrium</i>	0.04	0	0.01	Type 9	0.2	0.04				
λ_{EQ}		--	4.3	Type 10	0.1	0.04				
				Type 11	0	0.04				
Log-likelihood	-668.33	-702.57	-710.53			-677.65		-706.00		-715.77

Model	Level-<i>k</i> plus equilibrium			Mixture of Quantal Response Equilibrium types					
Specification	Individual specific precision (λ_i)	Type specific precision (λ_k)	Same precision (λ)	Individual specific precision (λ_i)		Type specific precision (λ_k)		Same precision (λ)	
<i>Random L0</i> ($\lambda = 0$)	0	--	--	<i>RL0</i> ($\lambda = 0$)				--	--
<i>Random L1</i>	0.60	0.98	0.99	$\lambda > 0$	1	QRE	0.80	λ	3.14
λ_{RL1}		8.54	8.71			λ_1	2.74		
<i>Random L2</i>	0.04	0	0			QRE	0.20		
λ_{RL2}		--	8.71			λ_2	9.63		
<i>Truthful L1</i>	0.14	0	0						
λ_{TL1}		--	8.71						
<i>Truthful L2</i>	0.01	0	0						
λ_{TL2}		--	8.71						
<i>Equilibrium</i>	0.21	0.02	0.01						
λ_{EQ}		29.84	8.71						
Log-likelihood	-569.30	-642.91	-644.12		-624.28		-684.81		-688.44

^bThis summary of the estimates for Goeree, Holt, and Palfrey pools their low- and high-value treatments for simplicity.