ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 75, No. 6 (November, 2007), 1721–1770

LEVEL-k AUCTIONS: CAN A NONEQUILIBRIUM MODEL OF STRATEGIC THINKING EXPLAIN THE WINNER'S CURSE AND OVERBIDDING IN PRIVATE-VALUE AUCTIONS?

VINCENT P. CRAWFORD
University of California, San Diego, La Jolla, CA 92093-0508, U.S.A.

NAGORE IRIBERRI
Universitat Pompeu Fabra, 08005 Barcelona, Spain

Econometrica, Vol. 75, No. 6 (November, 2007), 1721–1770

LEVEL-*k* AUCTIONS: CAN A NONEQUILIBRIUM MODEL OF STRATEGIC THINKING EXPLAIN THE WINNER'S CURSE AND OVERBIDDING IN PRIVATE-VALUE AUCTIONS?

BY VINCENT P. CRAWFORD AND NAGORE IRIBERRI[1]

The common curse of mankind, folly and ignorance, be thine in great revenue!
William Shakespeare, *Troilus and Cressida*

This paper proposes a structural nonequilibrium 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 nonequilibrium bidding in auction experiments. The level-k model generalizes many insights from equilibrium auction theory. It 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, nonequilibrium strategic thinking, behavioral game theory, experiments.

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 (MW) (1982); see the surveys by McAfee and McMillan (1987, Section X), Milgrom (1985, Section 4; 1987, Section 4), Wilson (1992, Section 4.2), and Klemperer (1999).

A central problem in this area is explaining the "winner's curse," the frequent tendency for bidders in common-value auctions to overbid, relative to

[1]Iriberri's work on this project began when she was affiliated with the University of California, San Diego. We are grateful to the National Science Foundation (Crawford), the Centro de Formación del Banco de España (Iriberri), and the Barcelona Economics Program of CREA (Iriberri) for research support; to Pierpaolo Battigalli, Yong Chao, Gary Charness, Olivier Compte, Ignacio Esponda, Erik Eyster, Drew Fudenberg, Charles Holt, Mark Isaac, Philippe Jehiel, John Kagel, Navin Kartik, Muriel Niederle, David Miller, Thomas Palfrey, Charles Plott, Matthew Rabin, Jose Gonzalo Rangel, Ricardo Serrano-Padial, Joel Sobel, Yixiao Sun, and, especially, Dan Levin for helpful discussions; and to Kagel for locating and providing the data from the experiments of Kagel and Levin (1986, 2002), Garvin and Kagel (1994), Kagel, Levin, and Harstad (1995), and Avery and Kagel (1997). A web appendix provides further analysis and detailed calculations.

equilibrium.[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 (KL) (1986), Kagel, Harstad, and Levin (1987), Lind and Plott (LP) (1991), Kagel, Levin, and Harstad (1995), and the papers surveyed in Kagel (1995, Section II) and KL (2002)). Finally, curse-like phenomena have been observed in nonauction 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 (HS) (1994), 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 (though not universal) tendency for subjects to bid higher than in the risk-neutral Bayesian equilibrium—though not usually to the point of making losses, on average, as in common-value auctions; see Cox, Smith, and Walker (1983, 1988); Goeree, Holt, and Palfrey (GHP) (2002), and the references cited therein.

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 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) formalized the intuition behind the curse in models in which "naive" bidders do not adjust their value estimates for the information revealed by winning, but otherwise follow equilibrium logic. Eyster and Rabin's (ER) (2002, 2005) notion of "cursed equilibrium" generalizes KL's and

[2]Some researchers use a more stringent definition: that the winner bids more than the expected value conditional on winning. Our weaker definition corresponds more closely to the deviations from equilibrium that are our main focus.

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

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 generalized 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 nonequilibrium 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 ample opportunity to learn from experience with analogous auctions,[6] but some auctions that have been studied using

[4]In Samuelson and Bazerman's (1985) Acquiring a Company experiments, 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- but not a more-informed player. It is unclear how to extend this interpretation of Samuelson and Bazerman's results 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 that he assumes all others share it—seems sensible. In related work, Esponda (2005) proposed a model in the spirit of self-confirming equilibrium (Fudenberg and Levine (1993)) to explain systematic deviations from equilibrium in games like Acquiring a Company. Jehiel and Koessler (2005) (see also Jehiel (2005)) proposed a general model of behavior in incomplete-information games in which players mentally bundle others' private-information types into analogy classes, which in a leading case reduces to fully cursed equilibrium. Like ER's notion, Esponda's and Jehiel and Koessler's notions are steady-state concepts meant to describe the outcome of a learning process.

[5]QRE 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. 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 exception of HS (2000). As LP (1991) noted, common-value auctions with risk aversion are not well understood theoretically.

[6]Such experience might justify fully-cursed equilibrium, for instance, by teaching bidders the trade-off 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' values, 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,

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.[7]

Progress via relaxing equilibrium requires a structural model that accurately describes initial responses to games.[8] In this paper we reconsider the winner's curse in common-value auctions and overbidding in independent-private-value auctions using nonequilibrium models of initial responses based on "level-k" thinking, introduced by Stahl and Wilson (1994, 1995) and Nagel (1995), and further developed and applied by Ho, Camerer, and Weigelt (1998), Costa-Gomes, Crawford, and Broseta (2001), Bosch-Domènech, Montalvo, Nagel, and Satorra (2002), Crawford (2003), Camerer, Ho, and Chong (CHC) (2004), Costa-Gomes and Crawford (2006), and Crawford and Iriberri (2007a). The level-k model has strong experimental support, which should allay the 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.

A level-k analysis 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

because once one realizes there may be a relationship to look for, there is no obvious reason to stop at intermediate levels of cursedness.

[7]Compare Fudenberg (2006): "...the fact that the amount of 'cursedness' typically declines as subjects become more experienced suggests that the curse, while real, is not an equilibrium phenomenon." It should eventually be possible to adapt the insights into cognition from analyses of initial responses to yield a deeper understanding of learning. Combining the two should then yield a clearer view of behavior in dynamic settings. Interesting evidence on learning in auctions is reported in Garvin and Kagel (1994), Kagel and Richard (2001), Neugebauer and Selten (2006), and Filiz-Ozbay and Ozbay (2007). Neugebauer and Selten's results for initial responses of subjects playing against random computer-simulated bidders include more underbidding than overbidding, and so suggest that some overbidding is a learned response, highly dependent on the feedback about the highest bid among other bidders.

[8]Maintaining 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; it also duplicates equilibrium in independent-private-value second-price auctions. By contrast, our approach dispenses with common knowledge of rationality (and of beliefs), but normally yields unique predictions.

empirical auction studies and nonauction experiments on strategic thinking, and thereby to bring a large body of auction evidence to bear on the issue of how best to model initial responses to games. Finally, it allows us to explore the issues that arise in extending level-k models to games of incomplete information and the robustness of standard auction theory's conclusions to failures of the equilibrium assumption.

A 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 Lk for $k > 0$ anchors its beliefs in a nonstrategic $L0$ type and adjusts them via thought experiments with iterated best responses: $L1$ best responds to $L0$, $L2$ to $L1$, et cetera. $L1$ and $L2$ have accurate models of the game and are rational; they depart from equilibrium only in basing their beliefs on simplified models of other players.[9] This yields a workable model of others' decisions while avoiding much of the cognitive complexity of equilibrium analysis.[10] In applications the population type distribution is usually translated from previous work or estimated from the current data set. The estimated distribution tends to be stable across games, with most of the weight on $L1$ and $L2$. Thus the anchoring $L0$ type exists mainly in the minds of higher types.

Even so, the specification of $L0$ 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 consider two alternative specifications,

[9] Charness and Levin (2005) conducted 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 mathematically the same as the one that underlies the curse. They find that the curse persists in their treatment, and conclude that their results favor explanations based on limited cognition in Bayesian updating or in understanding the problem rather than simplified models of others. Their results do 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 a level-k model can describe initial responses to environments, interactive or not, that pose cognitive difficulties isomorphic to those of predicting other players' strategic decisions. Dorsey and Razzolini (2003) reported experiments in which subjects made decisions in independent-private-value first-price auctions and lotteries that duplicate bidders' incentives in equilibrium. Their lotteries yield some overbidding, though less than their auctions, which suggests that overbidding is due in part to limited cognition.

[10] In 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 (2006) summarized the evidence for the level-k model and gave support for our assumptions that $L2$ best responds to an $L1$ without decision errors, unlike in Stahl and Wilson (1994, 1995), and to $L1$ alone rather than a mixture of $L1$ and $L0$, unlike *Worldly* in Stahl and Wilson (1995) and $L2$ in CHC (2004). We confine attention to $L0$, $L1$, and $L2$ because they well illustrate the model's potential to explain auction behavior and the evidence suggests that other types are comparatively rare.

both nonstrategic as is usual in level-k analyses. A *random* $L0$ 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 (2006). A *truthful* $L0$ bids the value that its own private information suggests, taken by itself, as in LP's (1991) naive model. We discuss these specifications of $L0$ in detail in Section 3. 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 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 their influences are altered by types' nonequilibrium beliefs. The pattern of types' deviations across first- and second-price common- and independent-private-value auctions determines whether a level-k model with a sensible type distribution can explain the systematic deviations from equilibrium such auctions often evoke.

Our analysis yields two main conclusions. First, many insights of equilibrium auction theory extend, suitably interpreted, to level-k auction theory.[11] Second, an empirically plausible level-k model can explain the winner's curse in common-value auctions and overbidding in independent-private-value auctions without the uniform value distributions used in most experiments.[12]

In common-value auctions, because random $L0$'s bids are independent of its signal, random $L1$ ignores the information revealed by winning, just as ER's

[11]To 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)).

[12]Gneezy (2005) reported experiments in which subjects play stylized common-value first- and second-price auctions with complete information. He finds that equilibrium predicts poorly and a level-k model like CHC's fits better than equilibrium in the second-price but not the first-price auction. (His first-price auctions yield results like those for the Traveler's Dilemma, whose structure is similar (see, for example, Goeree and Holt (2001)).) Gneezy's complete-information auctions and the Traveler's Dilemma raise significantly different behavioral issues than the auctions with diffuse private information considered here. Compte (2004) proposed an explanation of overbidding in both independent-private-value and common-value second-price auctions in which the key assumption is that bidders are overconfident in the accuracy of their own signals. In his model, the highest bidder is likely to be one whose error made him overoptimistic about his signal and so likely to overbid even in an independent-private-value auction. While such errors may make the model more realistic in applications, subjects in the experiments we study are told their signals with no error and no ambiguity of interpretation, so Compte's explanation is less plausible here.

fully-cursed equilibrium bidders do. In a second-price auction, the bidding trade-off is neutral and the lack of value adjustment makes random $L1$'s bids coincide with fully-cursed equilibrium bids, so that it normally overbids relative to equilibrium. In a first-price auction, random $L1$ differs from a fully-cursed-equilibrium bidder in using its nonequilibrium beliefs to evaluate a nonneutral bidding trade-off; this may make it bid higher or lower than fully-cursed equilibrium or coincide with it. In independent-private-value auctions (first- or second-price) with uniform values, random $L1$ coincides with equilibrium. Without uniformity, in general, random $L1$ may underbid, overbid, or coincide with equilibrium in first-price auctions, and it always coincides with equilibrium in second-price auctions.

In a first- or second-price auction, random $L1$'s bidding strategy is increasing in its signal. Thus in common-value auctions, random $L2$ adjusts its value estimate for the information revealed by winning. In a second-price auction, random $L2$ bids the expected value given its own signal, conditional on *just* winning. In this it follows the same logic as the equilibrium bidding strategy, but its beliefs do not anticipate winning if and only if it has the highest signal, which leads to a different adjustment. Value adjustment tends to make bidders' bids strategic substitutes, because winning against higher others' bids means others' signals are (stochastically) lower, which lowers the expected value conditional on winning. In a second-price auction, only value adjustment is relevant, so to the extent that random $L1$ overbids relative to equilibrium, random $L2$ underbids. In a first-price auction, the bidding trade-off may either reinforce or work against this tendency to underbid. In a first- or second-price independent-private-value auction, value adjustment is irrelevant. With uniform values, the bidding trade-off is neutral and random $L2$ coincides with equilibrium even in first-price auctions. With nonuniform values, random $L2$ coincides with equilibrium in second-price auctions, but it may underbid or overbid in first-price auctions.

In first- or second-price common-value auctions, truthful $L1$ tends to underbid relative to equilibrium or coincide with it. Truthful $L2$ tends to overbid or coincide with equilibrium. With uniform, independent private values, truthful $L1$ and $L2$ bids coincide with equilibrium. With nonuniform values, truthful $L1$ and $L2$ may underbid, overbid, or coincide with equilibrium in first-price auctions and they coincide with equilibrium in second-price auctions.

These bidding patterns allow a level-k model with an empirically plausible type distribution in which random $L1$ predominates, with lower frequencies of random $L2$, truthful $L1$ and $L2$, and an *Equilibrium* type that makes its equilibrium bid, to fit experimental data for common-value auctions better than equilibrium or cursed equilibrium, and to fit GHP's data for nonuniform independent-private-value auctions better than equilibrium or QRE. A level-k model has further advantages over cursed equilibrium in that it uses more general strategic principles to explain subjects' bidding behavior, with behavioral parameters linked to other bodies of evidence, and it may explain nonequilibrium bidding in some other independent-private-value auctions.

The rest of the paper is organized as follows. Section 2 introduces MW's (1982) general model with interdependent values and affiliated signals, and reviews the theories of equilibrium and cursed-equilibrium bidding. Section 3 discusses the specification of a level-k model for auctions and derives its general implications for random and truthful types. Section 4 compares equilibrium, cursed equilibrium, and the level-k model's implications in the leading examples that have been most often studied in auction experiments. Section 4 starts with the two common-value 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 (AK) (1997). It continues with second-price auctions in KL's example (for which ER (2002) but not ER (2005) discussed KL's results). Finally, since independent-private-value auctions are especially useful for separating cursed equilibrium from level-k decision rules, Section 4 analyzes GHP's (2002) design with discrete, slightly nonuniform values, in which level-k decision rules are separated, although weakly, from equilibrium. Section 5 compares the models econometrically in these four environments, using data on the initial responses of inexperienced subjects, which allow the cleanest tests of models of initial responses. 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 (2002) cursed equilibrium includes 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 is $V_i = u_i(S, X)$, where $u_i(\cdot, \cdot)$ is symmetric across i in that there is a function $u(\cdot, \cdot)$ on R^{M+N} such that for all i, $u_i(S, X) = u(S, X_i, \{X_j\}_{j\neq i})$, so that all bidders' valuations depend on S in the same manner and each bidder's valuation is a symmetric function of the other bidders' signals. Furthermore, 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 $V_i = u_i(X_i)$; the pure common-value model (used in KL (1986) and LP (1991)), in which $M = 1$ and $V_i = u_i(S)$; and an alternative common-value model (used in AK (1997)), in which $V_i = u_i(X) = \sum_{n=1}^{N} X_n$.

Because bidders are risk-neutral, if bidder i wins the auction and pays price p for the object, his payoff is $V_i - p$. For each i, Y, the highest signal among bidders other than i, has distribution function and density function, conditional on the realization x of X_i, $F_Y(y|x)$ 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: the expected value conditional on winning, $v(x, y) \equiv E[V_i|X_i = x, Y = y]$, and the unconditional expected value $r(x) = E[V_i|X_i = x]$. These functions are symmetric across i because $u_i(\cdot, \cdot)$ is symmetric across i.

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 2C'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 by winning because in a symmetric equilibrium the winner's signal must be more favorable than others' signals. 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)

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

where $1_{\{\cdot\}}$ is the indicator function and $\underline{x}$ is the infimum of the support of Y. 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[13]:

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

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)):

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

$a_*(x)$ 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, $v(x, x) \equiv x$, 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 $F_Y(y)/F_Y(x)$ and the first-price equilibrium bidding strategy becomes

$$a_*(x) = x - \int_{\underline{x}}^{x} \frac{F_Y(y)}{F_Y(x)}\, dy \equiv E[Y|Y < X]. \tag{4}$$

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

$$\begin{aligned} &\max_b E\big[(V_i - b_*(Y))1_{\{b_*(Y)<b\}}|X_i\big] \\ &\quad = \max_b \int_{\underline{x}}^{b_*^{-1}(b)} (v(x, y) - v(y, y)) f_Y(y|x)\, dy. \end{aligned} \tag{5}$$

[13]MW (1982, pp. 1107–1108) showed 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' nonequilibrium 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.

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

$$b_*(x) = v(x, x). \tag{6}$$

With independent private values, (6) becomes

$$b_*(x) = x \tag{7}$$

and the equilibrium $b_*(x)$ is a weakly dominant strategy in this case. 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 $f_Y(y|x)$. The common-value equilibrium in (6) is truthful like the independent-private-value equilibrium in (7), but the common-value equilibrium in (6) is not a weakly dominant strategy because optimal value adjustment depends on others' bidding strategies, as Section 3's level-k analysis shows more concretely.

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 the relationship between others' bids and signals. Instead they act as if they believe that with probability χ, ER's 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 equilibrium and $\chi = 1$ yields fully-cursed equilibrium, in which bidders assume there is no relationship 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.[14]

ER (2005, proof of Proposition 1, Proposition 5) simplified their analysis by showing that χ-cursed equilibrium is the same as equilibrium in a hypothetical

[14]The implicit assumption that a player gives too little weight to others' sophistication is often seen in other forms, for which it has considerable experimental support; see for example Weizsäcker (2003). As ER (2005, footnote 6) noted, cursed equilibrium allows certain kinds of differences in beliefs about others' type-contingent strategies.

"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 $u_i(S, X)$ is

$$V_i = u(S, X). \tag{8}$$

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)

$$\max_a \int_{\underline{x}}^{a_\chi^{-1}(a)} ((1-\chi)v(x, y) + \chi r(x) - a) f_Y(y|x)\, dy. \tag{9}$$

Just as (1) leads to (3), taking the partial derivative yields a differential equation whose solution determines the first-price χ-cursed-equilibrium bidding strategy:

$$a_\chi(x) = [(1-\chi)v(x, x) + \chi r(x)] - \int_{\underline{x}}^{x} \exp\left(-\int_y^x \frac{f_Y(t|t)}{F_Y(t|t)}\, dt\right) d[(1-\chi)v(y, y) + \chi r(y)]. \tag{10}$$

Like the first-price equilibrium bidding strategy $a_*(x)$, $a_\chi(x)$ 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' bids, he responds to the bidding trade-off just as an equilibrium bidder would. The effect of his 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 a cursed-equilibrium bidder overbids, 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)

$$\max_b \int_{\underline{x}}^{b_\chi^{-1}(b)} ((1-\chi)v(x,y) + \chi r(x) - (1-\chi)v(y,y) - \chi r(y)) f_Y(y|x)\,dy, \tag{11}$$

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

$$b_\chi(x) = (1-\chi)v(x,x) + \chi r(x). \tag{12}$$

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

From now on, we characterize the optimal bidding function only for the general model with interdependent values and affiliated signals; independent private values are a special case.

3. LEVEL-k MODELS

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, but it assumes that each bidder's behavior is drawn from a common distribution over a hierarchy of decision rules or *types*, in which $L1$ best responds to a nonstrategic anchoring type $L0$, $L2$ best responds to $L1$, et cetera. In this section, we derive types' implications in general; in Section 4, we specialize them to the examples used in the leading auction experiments.[15]

Bearing in mind that $L0$ represents only the starting point of a player's strategic thinking, as he processes the implications of the rules of the game before starting to think about others' responses to it, we consider two alternative specifications, each as naive as possible: A *random* $L0$ bids uniformly randomly, independent of its own private signal, over the entire range determined by the range of its signal and the value function $V_i = u_i(S, X)$.[16] A *truthful* $L0$ bids the

[15]Because any convex combination of monotonically increasing belief functions is monotonically increasing, hence invertible, which is all that is needed for our analysis, one could easily carry it out for CHC's cognitive hierarchy specification. Such an analysis would probably yield results close to ours (even allowing types higher than $L2$). We do not pursue this possibility because there is at least as much experimental support for our specification as CHC's (Costa-Gomes and Crawford (2006)) and our specification greatly simplifies characterizing types' implications.

[16]One can imagine more refined specifications of random $L0$, for example with bids uniformly distributed below its value instead of over the entire range of bids that are sensible for some

value its own signal suggests, taken by itself. We consider truthful as well as random types because it seems implausible to rule them out for all subjects a priori, but we report how the results change when truthful types are excluded so readers with stronger priors can draw their own conclusions.[17] We assume that a subject follows a given type, either random or truthful $L0, L1$, or $L2$ (footnote 10). Recall that "random" (or "truthful") $L1$ or $L2$ is shorthand for an $L1$ or $L2$ associated with a random (or truthful) $L0$; random or truthful $L1$ or $L2$ types need not be random or truthful themselves.

3A. *Random $L1$ and $L2$ Bidding Strategies in First- and Second-Price Auctions*

Random $L1$ assumes that other bidders are random $L0$, hence with bids independently and identically distributed (henceforth i.i.d.) uniformly over the range $[\underline{z}, \overline{z}]$ determined by the range of its private signal and the value function $V_i = u_i(S, X)$. Random $L1$ therefore believes that winning conveys no information about the value of the object, even with common values and affiliated

value. We avoid such refinements because $L0$ is meant to represent only the starting point for a player's analysis of others' bids, and it seems appropriate to follow the naive $L0$'s in most of the previous level-k literature, reserving more sophisticated thinking for $L1$ and $L2$. (Crawford and Iriberri (2007a) discussed this specification issue in detail. Complete-information models that adapt random $L0$ 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 Iriberri's (2007a) analysis of hide-and-seek games, where $L0$'s choice probabilities respond to the nonneutral framing of locations.) In the only other incomplete-information level-k models of which we are aware, CHC (2004, Section VI.A) used their closely related "cognitive hierarchy" model, with a random $L0$ defined like ours, 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. Brandts and Holt (1992, 1993) (see also Partow and Schotter (1993)) reported experiments with signaling games and showed that most of their sender and receiver subjects' initial responses are well described by their *naive thought process*, which is defined like our random $L1$. Our random $L1$ is also close to LP's (1991) private-value model. Ultimately the best specification of $L0$ is an empirical question and our random $L0$ allows a simple, coherent account of the data from auction experiments. 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.

[17] Our truthful $L0$ is equivalent to LP's (1991) naive model and reminiscent of the truthful sender type $W0$ in Crawford's (2003) level-k analysis of strategic deception via cheap talk. It also receives support in communication experiments (Crawford (1998), Cai and Wang (2006), and the papers cited there) and appears frequently in the informal literature on deception. Truthfulness is meaningful in auctions even though it has no natural meaning in most other settings for which level-k analyses have been conducted, but there is a nontrivial issue about how to define it. We take it to mean bidding $r(x)$ rather than x, so that subjects understand the strategic environment even if they lack sophistication in predicting others' responses to it. In independent-private-value auctions, and in KL's first- and second-price common-value auctions (Section 4A), $r(x) = x$, so this distinction is irrelevant. However, in settings like AK's common-value auctions, where the valuation is the sum of a player's own signal and his opponent's, it seems more sensible for even a naive subject to bid $r(x)$, which then equals his own signal plus the conditional expected value of the other's signal, than to bid just his own signal x.

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' uniformly random bids, Z, described by the distribution function $F_Z(z) = (\frac{z-\underline{z}}{\bar{z}-\underline{z}})^{N-1}$ and the density $f_Z(z) = (N-1)(\frac{z-\underline{z}}{\bar{z}-\underline{z}})^{N-2}\frac{1}{\bar{z}-\underline{z}}$. Note that these do not depend on the bidder's own signal x, which is uninformative about Z, or on the distribution of others' signals.

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

$$\text{(13)} \qquad \max_a E[(V_i - a)1_{\{Z<a\}}|X_i] \equiv \max_a \int_{\underline{z}}^{a} (r(x) - a) f_Z(z)\, dz \equiv \max_a (r(x) - a) F_Z(a).$$

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

$$\text{(14)} \qquad (r(x) - a) f_Z(a) - F_Z(a) = 0.$$

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 (13) and density and distribution function in (14) refer to random $L1$'s beliefs about the highest of $L0$ others' bids Z, 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 random $L1$ believes that winning conveys no information about the value of the object. Given the normal tendency for $r(x) > v(x, x)$, this tends to make random $L1$ overbid relative to equilibrium, just as a fully-cursed equilibrium bidder does. The second difference reflects random $L1$'s use of its nonequilibrium beliefs to evaluate the bidding trade-off between a higher bid's cost and increased probability of winning. Depending on the signal distribution, this difference may tend to either raise or lower random $L1$'s first-price bidding strategy relative to the equilibrium bidding strategy.

In a second-price auction, a random $L1$ bidder i's optimal bidding strategy solves

$$\text{(15)} \qquad \max_b E[(V_i - Z)1_{\{Z<b\}}|X_i] = \max_b \int_{\underline{z}}^{b} (r(x) - z) f_Z(z)\, dz.$$

Random $L1$'s second-price bidding strategy, $b_1^r(x)$, is characterized by the first-order condition

$$\text{(16)} \qquad (r(x) - b) f_Z(b) = 0 \quad \text{or, solving for } b, \quad b_1^r(x) = r(x).$$

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 random

$L1$'s nonequilibrium beliefs. But given random $L1$'s cursed value adjustment, truthful bidding is optimal, just as it is in an equilibrium analysis.[18] 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 $L1$ 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. Random $L1$'s bidding strategy therefore coincides with the second-price fully-cursed equilibrium bidding strategy in (12) 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 its weakly dominant strategy.

Unlike random $L1$, random $L2$ adjusts its value estimate for the information revealed by winning, because random $L1$'s bidding strategy is an increasing function of its private signal in either kind of auction.[19] We derive the optimal bids more generally, because the results will determine truthful $L1$'s and $L2$'s bidding strategies as well as random $L2$'s.

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)

$$\max_b E\big[(V_i - b_{k-1}(Y))1_{\{b_{k-1}(Y)<b\}}|X_i\big] = \max_b \int_{\underline{x}}^{b_{k-1}^{-1}(b)} (v(x,y) - b_{k-1}(y))f_Y(y|x)\,dy. \tag{17}$$

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

$$\big(v(x, b_{k-1}^{-1}(b)) - b\big) f_Y(b_{k-1}^{-1}(b)|x) \frac{\partial b_{k-1}^{-1}(b)}{\partial b} = 0 \quad \text{or} \quad v(x, b_{k-1}^{-1}(b)) - b = 0. \tag{18}$$

With independent private values, (18) reduces to the weakly dominant strategy in (7).

[18]Fully-cursed equilibrium and random $L1$ 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 $L1$ and fully-cursed equilibrium are not directly comparable, because random $L1$'s and equilibrium beliefs can differ considerably, depending on the specific distribution of the signals.

[19]This is easily verified from (14) for first-price auctions and (16) for second-price auctions.

Comparing the second-price level-k bidding strategy from (18) with the second-price equilibrium bidding strategy from (6) isolates the effect 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 the bidder has the highest signal, as an (symmetric) equilibrium bidder's do. A level-k 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 (18) 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 being equal.

Now 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)

$$\max_a E[(V_i - a)1_{\{a_{k-1}(Y)<a\}}|X_i] = \max_a \int_{\underline{x}}^{a_{k-1}^{-1}(a)} (v(x, y) - a) f_Y(y|x)\, dy. \tag{19}$$

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

$$\big(v(x, a_{k-1}^{-1}(a)) - a\big) f_Y(a_{k-1}^{-1}(a)|x) \frac{\partial a_{k-1}^{-1}(a)}{\partial a} - F_Y(a_{k-1}^{-1}(a)|x) = 0. \tag{20}$$

Comparing the first-price level-k bidding strategy determined by (20) 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, value adjustment interacts with the bidding trade-off, which, depending on the signal distribution and how the others' anticipated strategy $a_{k-1}(x)$ relates to the equilibrium strategy, may tend to either raise or lower the level-k bidding strategy relative to the equilibrium strategy. The web appendix (Crawford and Iriberri (2007b)) investigates this interaction in more detail, identifying the general principles that determine whether types overbid, underbid, or coincide with equilibrium here and in Section 4's examples.

Now consider how random $L2$'s first-price bidding strategy, $a_2^r(x)$, is determined by (20) with $a_1^{r^{-1}}(a)$ replacing $a_{k-1}^{-1}(a)$, hence by

$$\big(v(x, a_1^{r^{-1}}(a)) - a\big) f_Y(a_1^{r^{-1}}(a)|x) \frac{\partial a_1^{r^{-1}}(a)}{\partial a} - F_Y(a_1^{r^{-1}}(a)|x) = 0. \tag{21}$$

In a first-price auction, random $L2$, like random $L1$, deviates from equilibrium both in value adjustment and in using its nonequilibrium beliefs to evaluate the bidding trade-off. 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 random $L1$ overbids relative to equilibrium, because 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, random $L2$ believes that the curse is more severe than in equilibrium, and this tends to make it underbid, relative to equilibrium. Depending on the signal distribution and how random $L1$'s bidding strategy relates to the equilibrium strategy, the bidding trade-off may tend to raise or lower random $L2$'s bids relative to equilibrium or cursed equilibrium.

Random $L2$'s second-price bidding strategy, $b_2^r(x)$, is determined by (18) with $b_1^{r-1}(b)$ replacing $b_{k-1}^{-1}(b)$:

$$(22)\qquad \big(v(x, b_1^{r-1}(b)) - b\big) f_Y(b_1^{r-1}(b)|x)\frac{\partial b_1^{r-1}(b)}{\partial b} = 0 \quad \text{or}$$

$$b = v(x, b_1^{r-1}(b)).$$

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

3B. *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, truthful $L1$'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 (19). $a_1^t(x)$ is then determined by the first-order condition (20) with $a_0^{t-1}(a) \equiv r^{-1}(a)$ (because $a_0^t(x) \equiv r(x)$) replacing $a_{k-1}^{-1}(a)$:

$$(23)\qquad \big(v(x, r^{-1}(a)) - a\big) f_Y(r^{-1}(a)|x)\frac{\partial r^{-1}(a)}{\partial a} - F_Y(r^{-1}(a)|x) = 0.$$

Thus, in a first-price auction, truthful $L1$ deviates from equilibrium in its use of its nonequilibrium beliefs to evaluate the bidding trade-off, like random $L1$, but its different beliefs imply a different value adjustment.[20] Truthful $L0$

[20]Because 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 $v(x, y)$.

overbids relative to the first-price equilibrium bidding strategy, because it neither adjusts for the curse nor shades its bids. Hence truthful $L1$, which believes that to win it must bid higher than all others' truthful bids, not just higher than their equilibrium bids, believes that the curse is even more severe than in equilibrium. Thus the strategic substitutability of value adjustment tends to make truthful $L1$ underbid. But the bidding trade-off may again tend to raise or lower truthful $L1$'s bids relative to equilibrium.

In a second-price auction, a truthful $L1$ bidder's optimal bidding strategy, $b_1^t(x)$, solves a special case of the general monotonic problem (17) (for each x). Truthful $L1$'s second-price bidding strategy, $b_1^t(x)$, is then determined by (18) with $b_0^{t\,-1}(b) \equiv r^{-1}(b)$ replacing $b_{k-1}^{-1}(b)$:

$$(24) \qquad \big(v(x, r^{-1}(b)) - b\big) f_Y(r^{-1}(b)|x) \frac{\partial r^{-1}(b)}{\partial b} = 0 \quad \text{or} \quad b = v(x, r^{-1}(b)).$$

Thus, bidding is truthful as in the previous second-price analyses. Truthful $L0$ normally overbids relative to second-price equilibrium because it does not adjust for the curse, hence the strategic substitutability of value adjustment normally makes truthful $L1$ underbid.[21] In a common-value second-price auction, truthful $L1$'s bidding strategy is identical to random $L2$'s, because random $L1$ bids the expected value of the item based on its own signal, just as truthful $L0$ does.

In a first-price auction, truthful $L2$ expects other bidders to bid according to the monotonic bidding strategy $a_1^t(x)$, with inverse $a_1^{t\,-1}(a)$. Truthful $L2$'s first-price bidding strategy, $a_2^t(x)$, is then determined by problem (19) with $a_1^{t\,-1}(a)$ replacing $a_{k-1}^{-1}(a)$:

$$(25) \qquad \big(v(x, a_1^{t\,-1}(a)) - a\big) f_Y(a_1^{t\,-1}(a)|x) \frac{\partial a_1^{t\,-1}(a)}{\partial a} - F_Y(a_1^{t\,-1}(a)|x) = 0.$$

Thus, to the extent that truthful $L1$ underbids, value adjustment tends to make truthful $L2$ overbid. But the bidding trade-off may again raise or lower truthful $L2$'s bids relative to equilibrium.

In a second-price auction, truthful $L2$ expects other bidders to bid according to the monotonic bidding strategy $b_1^t(x)$, with inverse $b_1^{t\,-1}(b)$. Truthful $L2$'s second-price bidding strategy, $b_2^t(x)$, is again determined by (18), now with $b_1^{t\,-1}(b)$ replacing $b_{k-1}^{-1}(b)$:

$$(26) \qquad \big(v(x, b_1^{t\,-1}(b)) - b\big) f_Y(b_1^{t\,-1}(b)|x) \frac{\partial b_1^{t\,-1}(b)}{\partial b} = 0 \quad \text{or}$$
$$b = v(x, b_1^{t\,-1}(b)).$$

[21]In a second-price auction with independent private values, truthful $L0$'s (but not random $L0$'s) bids coincide with equilibrium when a player's signal reveals the actual value with certainty.

To the extent that truthful $L1$ underbids, value adjustment again makes 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 and one independent-private-value example. This section introduces the examples and their equilibrium, cursed equilibrium, and level-k bidding strategies to assess the level-k model's potential to explain behavior in the experiments and in preparation for Section 5's econometric analysis. Calculations are in the web appendix.

4A. *Kagel and Levin's, Avery and Kagel's, and Goeree, Holt, and Palfrey's Examples*

In the first example, used in KL's (1986) analyses of first-price auctions and in LP's (1991) follow-up experiments, $N \geq 3$, $V_i = u_i(S, X) = S$, S is uniformly distributed on a subset of the real line $[\underline{s}, \overline{s}]$, 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 $\overline{s}$. 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

$$v(x, y) = \begin{cases} x - \dfrac{a}{2} + \dfrac{a}{N} - \dfrac{x-y}{N}, \\ \quad x - a \leq y \leq x, \\ \\ y - \dfrac{a}{2} + \dfrac{a}{N} - \dfrac{\left(\dfrac{y-x}{a}\right)^{N-1}}{\left[1 - \left(\dfrac{y-x}{a}\right)^{N-1}\right]}\left(\dfrac{N-1}{N}\right)(x + a - y), \\ \quad x < y \leq x + a. \end{cases} \tag{27}$$

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 (1997) analysis of second-price auctions, $V_i = u_i(S, X) = \sum_{i=1}^{N} X_i$ and X_i is i.i.d. uniformly distributed on the interval $[\underline{x}, \overline{x}]$. Thus, in general, $r(x) \equiv E[\sum_{k=1}^{N} X_k | X_i = x] = x + (N - 1)\frac{\overline{x}+\underline{x}}{2}$, $v(x, y) = x + y\frac{N}{2} + \frac{(N-2)}{2}\underline{x}$, and $v(x, x) = x + x\frac{N}{2} + \frac{N-2}{2}\underline{x} > (<)\ r(x)$ if and only if $x > (<)\ \frac{(N-1)\overline{x}+\underline{x}}{N}$, so that $v(x, x) > r(x)$ 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.[22] When $N = 2$ and $[\underline{x}, \overline{x}] = [1, 4]$, as in AK's experiments, $r(x) = x + \frac{5}{2}$ and $v(x, x) = 2x$, so that $r(x) < v(x, x)$ when $x > \frac{5}{2}$ and $r(x) > v(x, x)$ when $x < \frac{5}{2}$.

In the third example, used in GHP's (2002) analysis of first-price independent-private-value auctions, $N = 2, V_i = u_i(S, X) = X_i$, and there are two treatments, each with bids restricted to integer values and discrete, slightly nonuniform (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.

We now describe the relationships among equilibrium, cursed-equilibrium, and random and truthful $L1$ and $L2$ bidding strategies in the examples. Table I summarizes the conclusions, first in general and then in KL's and AK's examples. The conclusions follow fairly simply from the facts that only the bidding trade-off (as influenced by equilibrium, cursed equilibrium, or level-k beliefs) matters in first-price independent-private-value auctions, that only value adjustment (as influenced by the various beliefs) matters in second-price common-value auctions, and that the two effects combine in straightforward ways in first-price common-value auctions.

4B. *Equilibrium and Cursed Equilibrium versus Level-k Models in Second-Price Auctions*

In a second-price auction with independent private values, random and truthful $L1$ and $L2$ bid truthfully, as in equilibrium and cursed equilibrium, because they follow weakly dominant strategies when they exist. Thus neither level-k model can explain nonequilibrium bidding (except via random $L0$, which we estimate to have 0 or, in one case, low frequency).

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

In a second-price auction with common values, because random $L1$ bids the value its own signal suggests, like truthful $L0$, truthful $L1$ coincides with random $L2$.[23] We have not derived a closed-form solution for truthful $L2$ in the

[22]This 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.

[23]Although in independent-private-value auctions, random Lk types are equivalent to the analogous truthful Lk types when the distribution of private signals is unconditionally uniform, in

TABLE I

TYPES' BIDDING STRATEGIES[a,b]

Auction/Type	Equilibrium	χ-Cursed Equilibrium	Random $L1$	Random $L2$	Truthful $L1$	Truthful $L2$
2nd-price i.p.v.	x	x	$b_1^r(x) = x$	$b_2^r(x) = x$	$b_1^t(x)$ from (24) with $v(x, \cdot) \equiv x$	$b_2^t(x)$ from (26) with $v(x, \cdot) \equiv x$
2nd-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 (22): $b = v(x, b_1^{r-1}(b))$	$b_1^t(x)$ from (24): $b = v(x, r^{-1}(b))$	$b_2^t(x)$ from (26): $b = v(x, b_1^{t-1}(b))$
2nd-price c.v.: KL	$x - \frac{a}{2} + \frac{a}{N}$	$x - (1-\chi)a\frac{N-2}{2N}$	x	$x - \frac{a}{2}(\frac{N-2}{N-1})$	$x - \frac{a}{2}(\frac{N-2}{N-1})$	No closed-form solution
2nd-price c.v.: AK	$2x$	$\chi(x + \frac{5}{2}) + (1-\chi)2x$	$x + \frac{5}{2}$	3.5 if $x \leq 2.5$; 6.5 if $x > 2.5$	3.5 if $x \leq 2.5$; 6.5 if $x > 2.5$	No closed-form solution
1st-price i.p.v.	$a_*(x)$ from (4)	$a_*(x)$ from (4)	$a_1^r(x)$ from (14)	$a_2^r(x)$ from (21) with $v(x, \cdot) \equiv x$	$a_1^t(x)$ from (23) with $v(x, \cdot) \equiv x$	$a_2^t(x)$ from (25) with $v(x) = x$
1st-price c.v.	$a_*(x)$ from (3)	$a_\chi(x)$ from (10)	$a_1^r(x)$ from (14)	$a_2^r(x)$ from (21)	$a_1^t(x)$ from (23)	$a_2^t(x)$ from (25)
1st-price c.v.: KL	$x - \frac{a}{2} + \frac{a}{N+1}$ $\times \exp(-\frac{N(x - \underline{x} - \frac{a}{2})}{a})$	$[\chi x + (1-\chi)(x - \frac{a}{2} + \frac{a}{N})$ $- \frac{a}{N}] + \frac{a}{N+1}\exp(-\frac{N(x - \underline{x} - \frac{a}{2})}{a})$	$x - \frac{a}{N}$	$x - \frac{a}{2}$	$x - \frac{a}{2}$	$x - \frac{a}{2}$

[a] If there is no general closed-form expression, Table I refers to the equation in the text that determines the bidding strategy.

[b] The abbreviation i.p.v. denotes independent private value and c.v. denotes common value.

examples, but computations show that in KL's example truthful $L2$ overbids by more than a fully cursed-equilibrium bidder, to an extent that increases with N, and 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. With common values, a level-k model has the potential to improve upon cursed equilibrium, but this depends on whether an empirically plausible mixture of level-k types gives a better account of subjects' heterogeneous bidding behavior than a plausible mixture of cursed types.

4C. *Equilibrium and Cursed Equilibrium versus Level-k Models in First-Price Auctions*

In a first-price auction with independent private values, in general the bidding trade-off may tend to make a random or truthful $L1$ or $L2$ either underbid or overbid, depending on the value distribution. Most independent-private-value experiments used values uniformly i.i.d. on $[\underline{x}, \overline{x}]$. 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 $c(x - \underline{x}) + \underline{x}$, as long as $0 < c \leq 1$. Random or truthful $L1$, and therefore random or truthful $L2$, then coincide with equilibrium, and this limits the potential for a level-k model to improve upon an equilibrium explanation of overbidding.[24]

For nonuniform value distributions, a level-k model may be able to explain nonequilibrium bidding. In GHP's (2002) independent-private-value designs, random $L1$ or $L2$ coincides with equilibrium except for the highest valuation in the high-value treatment, where random $L1$ slightly overbids and random $L2$ underbids (see the web appendix). Truthful $L1$ underbids in the low-value and overbids in the high-value treatment, and truthful $L2$ underbids in both.[25]

In a first-price auction with common values, in KL's example, when $N = 2$, equilibrium and fully-cursed equilibrium bids coincide and random $L1$ bids

common-value auctions, random and the analogous truthful types are not equivalent in general, because they differ in value adjustment.

[24]Some potential for improvement remains because the costs of deviations differ slightly for *Equilibrium* and random $L1$, et cetera, so they are weakly separated. For low values and low precision, underbidding is less costly for *Equilibrium* than for random $L1$, while overbidding is more costly for *Equilibrium* than for random $L1$. As precision increases, this asymmetry between under- and overbidding disappears except for very low values, and both under- and overbidding are costlier for *Equilibrium*. For high values, both under- and overbidding are costlier for *Equilibrium* than for random $L1$, so the $L1$ probability distribution of decisions has thicker tails. Differences in deviation costs sometimes separate types in other treatments (Section 5).

[25]In GHP we define random $L0$ with equal probabilities for subjects' possible values in each treatment. Random and truthful specifications do not coincide in GHP's design, even though the ex ante values and random $L0$ are uniformly distributed, because the values are discrete and unevenly spaced, and integer bids between the values are allowed.

are slightly lower than but approximately coincide with them; random $L2$ bids approximately coincide with equilibrium or fully-cursed equilibrium.[26] When $N > 2$, random $L1$ bids approximately coincide with fully-cursed equilibrium bids, but both overbid relative to equilibrium, by an amount that increases with N and a; random $L2$ bids approximately coincide with equilibrium, but underbid relative to fully-cursed equilibrium, by an amount that increases with N and a. Value adjustment and the bidding trade-off offset each other for random $L2$ and truthful $L1$, which approximately coincides with equilibrium. Truthful $L2$ approximately coincides with equilibrium because truthful $L1$ does.

To sum up for first-price auctions, with uniform independent private values, level-k types coincide with equilibrium and cursed equilibrium, but with nonuniform value distributions as in GHP (2002), random (or truthful) $L1$ bids coincide with (or fall below) equilibrium bids in the low-value treatment and exceed equilibrium bids in the high-value treatment; random (or truthful) $L2$ bids coincide with (or fall below) equilibrium bids for both treatments. A level-k model is then weakly separated from equilibrium and cursed equilibrium, and may be able to explain nonequilibrium bidding. With common values, a level-k model again has the potential to improve upon cursed equilibrium.

5. COMPARING THE MODELS ECONOMETRICALLY

All of the models compared here depend on behavioral parameters: logit error precisions for all of them, plus population type frequencies for level-k models or cursedness parameters for cursed-equilibrium models. This section uses existing data from auction experiments to estimate the models econometrically and compare their abilities to account for observed behavior in the experiments. Our goal in the econometrics 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. We estimate treatment by treatment: Because our main purpose is model evaluation and the treatments have widely differing subject populations and experimental conditions, we have not tried to pool them.

Table II summarizes the data we use. Because learning can lead even unsophisticated subjects to equilibrium, strategic thinking appears most clearly before subjects have seen others' responses. We therefore (unlike ER) use data only from inexperienced subjects and (instead of pooling data from all periods and usually all subjects as ER did) we focus on individual subjects' initial responses, interpreted as the first five periods (in which a subject typically had five different realizations of his private signal) to compensate for small sample size.

[26]"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. We follow them in this, for cursed equilibrium as well as equilibrium.

TABLE II

DATA SOURCES AND EXPERIMENTAL DESIGNS

G (treatment)	Auction Type	$u(S, X)$	Signals	n (Sample Size)	Treatment Variables
1. KL first-price	First-price common value	$u(S, X) = S$	$X \mid S \sim U[s - a/2, s + a/2]$	51	a (dispersion), N (number of bidders), limits of s
2. KL second-price	Second-price common value	$u(S, X) = S$	$X \mid S \sim U[s - a/2, s + a/2]$	28	a (dispersion)
3. AK second-price	Second-price common value	$u(S, X) = X_1 + X_2$	$X \sim U[\underline{x}, \overline{x}] = [1, 4]$	23	No variation, $N = 2$
4. GHP	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, $N = 2$

Given these choices, we maximize comparability with ER's analysis of KL's (1986) first-price and AK's (1997) second-price data. KL, however, had only experienced subjects (who had participated in at least one prior auction session), and while AK had some inexperienced subjects, ER's analysis focused on their experienced subjects. For common-value second-price auctions, we therefore use AK's data for inexperienced subjects and the unpublished data for inexperienced subjects in the second-price version of KL's design mentioned in the Appendix to Kagel, Levin, and Harstad (1995) as reprinted in KL (2002, Chapter 4), referring to the latter as the KL second-price data. For common-value first-price auctions, we use the data for inexperienced subjects in KL's design from Garvin and Kagel (1994) (reprinted in KL (2002, Chapter 10)), referring to them as the KL first-price data.[27]

Finally, because cursed equilibrium coincides with equilibrium in independent-private-value auctions, they are particularly important in assessing the level-k model, but with independent private values, level-k types 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 4C).[28] We therefore use GHP's (2002) data from first-price independent-private-value auctions with discrete nonuniform values, which weakly separate level-k types from equilibrium.[29]

Our econometric specification follows the mixture-of-types models of Stahl and Wilson (1994, 1995), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004), Costa-Gomes and Crawford (2006), and Crawford and Iriberri (2007a). Level-k and cursed types, *Equilibrium*, and QRE types are all assumed to make logistic errors as described below. (Random $L0$ directly specifies a uniform distribution of decisions and so has no precision parameter.)

[27]Other common-value experiments whose data would enrich our analysis include LP's (1991) and HS's (2000), but despite those authors' generous efforts, their data are unavailable.

[28]This coincidence extends even to Kagel and Levin's (1993) uniform independent-private-value third-price auctions.

[29]Goeree and Holt (2001, Section III) reported similar results for a closely related design, which we do not consider here (although their data are available). In Palfrey's (1985) and Chen and Plott's (1998) independent-private-value designs, level-k types also deviate from equilibrium, but despite their generous efforts, their data are unavailable. We define payoffs as payments for performance, omitting show-up fees, and express them in 1989 dollars. Following GHP, to avoid distorting the estimates, we edited a small number of "crazy" bids (6 in AK, 11 in KL first-price, 4 in KL second-price, and 12 in GHP)—3.6% of the sample—replacing bids above the highest (below the lowest) rationalizable bid with the highest (lowest) such bid. In the subject-by-subjects estimation, this editing affects only the edited subjects' type estimates. In the type-specific and constant-precision estimates, editing can potentially affect all subjects' estimates via the estimated precisions, but in these cases, *not* editing the data, and so requiring the error structure to "explain" crazy bids along with normal decision noise, seems more likely to distort the estimates than editing them, especially for inexperienced subjects. The alternative—throwing these subjects' data away—would yield similar estimates, but seems arbitrary and would discard data from subjects most of whose choices were sensible.

For our level-k plus equilibrium models, we allow random $L0$ and both random and truthful $L1$ and $L2$ types as well as *Equilibrium*, each with its own beliefs (Section 3). In footnotes 37, 39, and 40, we report how the results change with truthful types excluded; the details of those estimates are provided in Section C of the web appendix.[30] The most general specifications allow subject-specific precisions, but we also consider type-specific and constant precisions.

For our cursed-equilibrium models in the common-value treatments, we also allow random $L0$ to avoid biasing the comparisons. In the most general specification, we allow subject-specific precisions and levels of ER's cursedness parameter χ, as in their analysis of AK's data, but we also consider models with "cursed types," both with type-specific precision and constant precision. In the former case, for computational tractability, we constrain χ to multiples of 0.1 in $[0, 1]$. In the latter cases we constrain χ to a number of estimated values in $[0, 1]$ equal to the number of types in the analogous level-k model. Either way, unlike ER, we restrict χ to $[0, 1]$.[31] Each of our cursed-equilibrium models allows $\chi = 0$ and so nests equilibrium, which is important for a fair comparison of cursed-equilibrium and level-k models. We have more confidence in the cursed types $\chi = 0$ or 1 because their theoretical rationales are stronger than for intermediate values of χ (footnote 6), but estimates of models allowing intermediate values are useful diagnostics.

For GHP's first-price independent-private-value treatments, where cursed equilibrium coincides with equilibrium, we replace cursed equilibrium with a QRE model like the one GHP favor. Random $L0$ is implicitly included as a QRE type with 0 precision. In the most general specification, we again allow subject-specific precisions, but we also consider models with QRE types, both with type-specific and constant precision. We again constrain the number of types to that of the analogous level-k model.[32]

The formal discussion that follows covers all three models and all three error structures, with $k = 1, 2, \ldots, K$ indexing level-k (or *Equilibrium*) types, cursed

[30]We omit truthful $L0$ in the econometric analysis because truthful bidding is very rare for the first-price treatments (6/255 observations in KL and 6/400 in GHP, with no subject making more than two truthful bids) and because there is no way to assign beliefs that make 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.

[31]Unlike level-k models, cursed equilibrium can accommodate heterogeneous bidding behavior only via cursed types or subject-specific cursedness parameters. ER (2005, Table II) allowed χ to take any value and reported many estimates for AK's inexperienced subjects outside $[0, 1]$, contradicting χ's interpretation as a probability. This problem would also arise in unconstrained estimates for KL's examples, where the below-equilibrium or above-signal bids sometimes observed correspond to $\chi < 0$ or $\chi > 1$. Level-k types often explain such bids better than cursed types with $\chi = 0$ or 1, particularly in second-price common-value auctions like AK's and KL's (Table I).

[32]We depart from GHP by ruling out nonneutral risk preferences and payoffs for "joy of winning," in keeping with our goal of learning whether a level-k model can explain auction behavior without such augmentations.

types, or QRE types. Index Table II's treatments (first- or second-price) g (for "games") $= 1, 2, 3, 4$. Each type k implies a bidding strategy in game g, denoted $c_k^g(x)$; c_{it}^g denotes subject i's observed bid in game g at time t. We assume that a subject of type k normally follows $c_k^g(x)$, but subject to logistic errors of precision λ, assumed independent across the five periods in which he plays. Write his expected payoff for bid c given signal x with type k's beliefs $S_k^g(c|x)$ (formally defined in Section 2 or 3). The probability of observing bid c within the range of possible bids $[\underline{c}, \overline{c}]$ for type k is then

$$\text{(28)} \qquad \Pr(c|k, x, g, \lambda) = \frac{\exp(\lambda S_k^g(c|x))}{\int_{\underline{c}}^{\overline{c}} \exp(\lambda S_k^g(e|x))\, de}.$$

As usual, this implies that the costlier an error is ex ante, given type k's beliefs, the lower the player's probability of making it, with the cost sensitivity tuned by the precision λ. The player's bids approach uniform randomness as $\lambda \to 0$ or the error-free bid $c_k^g(x)$ as $\lambda \to \infty$.

The matrix $\Lambda \equiv [\lambda_{ik}]$ gives precision indexed by subject i and type k. Subject-specific precisions do not restrict how λ_{ik} varies with i and k. Type-specific precisions restrict λ_{ik} to be independent of i for any given k. Constant precisions restrict λ_{ik} to be independent of i and k.

With errors independent, conditional on type, the likelihood of observing the five-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_{ik} in game g is

$$\text{(29)} \qquad L_k(c_i^g|k, x, g, \lambda_{ik}) = \prod_{t=1}^{5} \Pr(c_{it}^g|k, x, g, \lambda_{ik}).$$

Let π_k denote the proportion of type k in the population, with $\sum_k \pi_k = 1$. The likelihood of observing c_i^g unconditional on type is then

$$\text{(30)} \qquad \sum_{k=1}^{K} \pi_k L_k(c_i^g|k, x, g, \lambda_{ik}) = \sum_{k=1}^{K} \pi_k \prod_{t=1}^{5} \Pr(c_{it}^g|k, x, g, \lambda_{ik}).$$

Given (28), because the payoff function is quasiconcave and the logit term increases with payoff, the likelihood treats a bid as stronger evidence for a type the closer it is to the type's bid or the better the deviations are explained given its beliefs. In most cases, types' bids differ and the first factor is more important. Although some types' bids always or almost always coincide, even they are usually weakly separated by differences in the deviation costs implied by their beliefs.

Indexing treatment g's subjects $i = 1, 2, \ldots, N_g$ and letting $c^g = (c_1^g, c_2^g, \ldots, c_{N_g}^g)$, from (30) we can now write the models' likelihood (L) and log-likelihood

(LL) functions for treatment g:

$$(31)\qquad L(\pi, \Lambda|c^g) = \prod_{i=1}^{N_g} \sum_{k=1}^{K} \pi_k L_k(c_i^g|k, x, g, \lambda_{ik}) \quad \text{and}$$

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

Tables IIIa–c summarize treatment-by-treatment parameter estimates and likelihoods for our level-k and cursed-equilibrium models for KL first- and second-price and AK second-price. Table IIId summarizes parameter estimates and likelihoods for our level-k and QRE models for GHP first-price. Standard errors except for the models with subject-specific precisions (in the web appendix (Crawford and Iriberri (2007b)) Sections D and E as explained below) are in parentheses.[33] Level-k types that are not separated from other types in a treatment are listed with equivalences indicated by a tilde ($\sim$).[34]

In each treatment, likelihood-ratio tests for the level-k plus equilibrium models, for which the alternative error structures are nested, strongly reject constant or type-specific error precisions (p-values < 0.0015). The Bayesian information criterion (BIC), which adjusts the likelihood to penalize models with more parameters without requiring that the models be nested, also favors models with subject-specific precisions, except in GHP, where it favors constant precisions for the level-k model and type-specific precisions for the QRE model.[35] For the cursed-equilibrium models, for which the error structures are

[33]Standard errors of the parameters in the models with type-specific or constant precision were obtained using a jackknife procedure. In each run of the jackknife, the model was reestimated with one subject excluded. The standard deviations of the parameter estimates across these runs are the estimated standard errors of the parameters. In each run, we initialized the computations at the estimates with all subjects included. Although in theory this choice can bias the estimated standard errors downward, initializing randomly yielded estimates only slightly higher. Standard errors of the precisions in the models with subject-specific precisions were also obtained using a jackknife procedure. In each run, the model was reestimated with one period excluded, with each subject's estimated type fixed at its value with all periods included. The standard deviations of the parameter estimates across these runs are the estimated standard errors of the subject-specific precisions.

[34]In KL first-price, random $L2$ and truthful $L1$ are separated from each other and from *Equilibrium* only by deviation costs (due to their different beliefs), and truthful $L2$ is not separated from *equilibrium* even by deviation costs. In the second-price auctions, truthful $L1$ and random $L2$ are not separated even by deviation costs (because their beliefs are the same), and neither are truthful $L2$ and random $L3$. In GHP first-price, random $L1$ and *equilibrium* are separated only by bids for $v = 12$ in the high-value treatment (web appendix) and by deviation costs for other values; random $L2$ and *equilibrium* are separated only in the high-value treatment and only by deviation costs. For simplicity, Table IIId pools the results for GHP's low- and high-value treatments.

[35]Here and below, the Akaike information criterion, which makes an adjustment similar to the BIC but requires that the models be nested, always orders nested models in the same way

TABLE IIIa

MODELS AND ESTIMATES FOR KAGEL AND LEVIN FIRST-PRICE

	Level-k Plus Equilibrium					Cursed Equilibrium								
Specification	Subject-Specific Precision (λ_i)	Type-Specific Precision (λ_k)		Constant Precision (λ)		Subject-Specific Precision (λ_i) and Fixed Cursedness Types ($\chi = (1, 0.9, \ldots, 0)$)			Type-Specific Precision (λ_k)			Constant Precision (λ)		
	$\hat{\pi}_k$	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\pi}_k$	$\hat{\lambda}$	Types	χ	$\hat{\pi}_k$	$\hat{\chi}_k$	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\chi}_k$	$\hat{\pi}_k$	$\hat{\lambda}$
Random $L0$	0.04	0	—	0	—	Random $L0$	—	0.06	—	0	—	—	0	—
		(0.00)		(0.00)						(0.00)			(0.00)	
Random $L1$	0.59	0.35	1	0.49	1.62	Type 1	1	0.47	0.99	0.83	0.60	1	0.50	0.68
		(0.04)	(0.14)	(0.02)	(0.11)				(0.03)	(0.01)	(0.02)	(0.00)	(0.02)	(0.02)
Random $L2$	0.04	0.03	280.90	0	1.62	Type 2	0.90	0.02	0.78	0.06	46.20	0	0.50	0.68
		(0.00)	(299.12)	(0.01)	(0.11)				(0.01)	(0.01)	(77.91)	(0.00)	(0.02)	(0.02)
Truthful $L1$	0.18	0.54	1.21	0.29	1.62	Type 3	0.80	0.08	0	0.11	14.74			
		(0.04)	(0.07)	(0.02)	(0.11)				(0.00)	(0.01)	(0.71)			
Truthful $L2$	~Eq.	~Eq.	~Eq.	~Eq.	~Eq.	Type 4	0.70	0.06						
Equilibrium	0.16	0.08	11.09	0.22	1.62	Type 5	0.60	0						
		(0.01)	(0.58)	(0.02)	(0.11)									
						Type 6	0.50	0						
						Type 7	0.40	0.04						
						Type 8	0.30	0.04						
						Type 9	0.20	0.04						
						Type 10	0.10	0						
						Type 11	0	0.20						
Log-likelihood	−1,658.30	−1,739.60		−1,753.54				−1,635.96		−1,736.62			−1,762.24	
BIC	−1,724.48	−1,749.23		−1,759.56				−1,710.56		−1,747.45			−1,768.26	

TABLE IIIb

MODELS AND ESTIMATES FOR KAGEL AND LEVIN SECOND-PRICE

	Level-k Plus Equilibrium					Cursed Equilibrium								
Specification	Subject-Specific Precision (λ_i)	Type-Specific Precision (λ_k)		Constant Precision (λ)		Subject-Specific Precision (λ_i) and Fixed Cursedness Types ($\chi = (1, 0.9, \ldots, 0)$)			Type-Specific Precision (λ_k)			Constant Precision (λ)		
	$\hat{\pi}_k$	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\pi}_k$	$\hat{\lambda}$	Types	χ	$\hat{\pi}_k$	$\hat{\chi}_k$	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\chi}_k$	$\hat{\pi}_k$	$\hat{\lambda}$
Random $L0$	0	0	—	0	—	Random $L0$	—	0.18	—	0.43	0	—	0	—
		(0.00)		(0.00)						(0.02)	(0.00)		(0.00)	
Random $L1$	0.21	0.10	95.84	0.62	8.91	Type 1	1	0.18	0.86	0.27	8.89	0.79	0.43	2.95
		(0.10)	(2.30)	(0.02)	(0.39)				(0.02)	(0.03)	(1.51)	(0.03)	(0.03)	(0.40)
Random $L2$	0.21	0.27	2.50	0.11	8.91	Type 2	0.9	0.11	0.18	0.30	5.35	0.33	0.15	2.95
		(0.02)	(0.58)	(0.01)	(0.39)				(0.01)	(0.04)	(0.30)	(0.09)	(0.04)	(0.40)
Truthful $L1$	$\sim R.L2$	$\sim R.L2$	$\sim R.L2$	$\sim R.L2$	$\sim R.L2$	Type 3	0.8	0.04				0	0.42	2.95
												(0.00)	(0.02)	(0.40)
Truthful $L2$	0.32	0.33	6.10	0.27	8.91	Type 4	0.7	0						
		(0.02)	(0.23)	(0.02)	(0.39)									
Equilibrium	0.25	0.30	49.76	0	8.91	Type 5	0.6	0.07						
		(0.02)	(3.17)	(0.00)	(0.39)									
						Type 6	0.5	0.04						
						Type 7	0.4	0.04						
						Type 8	0.3	0						
						Type 9	0.2	0.11						
						Type 10	0.1	0.07						
						Type 11	0	0.18						
Log-likelihood	−918.26	−967.80		−973.81				−950.91		−987.48			−995.59	
BIC	−952.60	−976.39		−979.17				−992.76		−997.14			−1,003.11	

TABLE IIIc

MODELS AND ESTIMATES FOR AVERY AND KAGEL SECOND-PRICE

	Level-k Plus Equilibrium					Cursed Equilibrium								
Specification	Subject-Specific Precision (λ_i)	Type-Specific Precision (λ_k)		Constant Precision (λ)		Subject-Specific Precision (λ_i) and Fixed Cursedness Types ($\chi = (1, 0.9, \ldots, 0)$)			Type-Specific Precision (λ_k)			Constant Precision (λ)		
	$\hat{\pi}_k$	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\pi}_k$	$\hat{\lambda}$	Types	χ	$\hat{\pi}_k$	$\hat{\chi}_k$	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\chi}_k$	$\hat{\pi}_k$	$\hat{\lambda}$
Random $L0$	0	0	—	0	—	Random $L0$	—	0.13	—	0	—	—	0	—
		(0.00)		(0.00)						(0.00)			(0.00)	
Random $L1$	0.65	0.56	12.75	0.94	4.3	Type 1	1	0.43	1	0.37	9.67	0.8	1	2.77
		(0.03)	(0.79)	(0.01)	(0.27)				(0.00)	(0.03)	(0.69)	(0.03)	(0.00)	(0.17)
Random $L2$	0.09	0	—	0.06	4.3	Type 2	0.9	0	0.73	0.08	161.45			
		(0.00)		(0.01)	(0.27)				(0.05)	(0.01)	(42.60)			
Truthful $L1$	$\sim R. L2$	$\sim R. L2$	$\sim R. L2$	$\sim R. L2$	$\sim R. L2$	Type 3	0.8	0	0.63	0.55	1.33			
									(0.06)	(0.03)	(0.15)			
Truthful $L2$	0.22	0.05	633.01	0	4.3	Type 4	0.7	0.13						
		(0.01)	(3.54)	(0.00)	(0.27)									
Equilibrium	0.04	0.39	0.63	0	4.3	Type 5	0.6	0.04						
		(0.03)	(0.20)	(0.00)	(0.27)									
						Type 6	0.5	0.09						
						Type 7	0.4	0.04						
						Type 8	0.3	0						
						Type 9	0.2	0.04						
						Type 10	0.1	0.04						
						Type 11	0	0.04						
Log-likelihood	−668.23	−702.44		−710.53				−677.65		−706.00			−715.77	
BIC	−696.05	−710.69		−715.68				−712.68		−715.27			−722.98	

TABLE IIId

MODELS AND ESTIMATES FOR GOEREE, HOLT, AND PALFREY FIRST-PRICE[a]

	Level-k Plus Equilibrium					QRE					
Specification	Subject-Specific Precision (λ_i)	Type-Specific Precision (λ_k)		Constant Precision (λ)		Subject-Specific Precision (λ_i)		Type-Specific Precision (λ_k)		Constant Precision (λ)	
	$\hat{\pi}_k$	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\pi}_k$	$\hat{\lambda}$	Types	$\hat{\pi}_k$	$\hat{\lambda}_k$	$\hat{\pi}_k$	$\hat{\lambda}$	$\hat{\pi}_k$
Random $L0$	0	0 (0.00)	—	0 (0.00)	—	Random $L0$	0	—	0 (0.00)	—	0 (0.00)
Random $L1$	0.65	0.98 (0.00)	8.54 (0.09)	0.99 (0.00)	8.71 (0.09)	$\hat{\lambda} > 0$	1	2.74 (0.02)	0.80 (0.01)	3.14 (0.03)	1 (0.00)
Random $L2$	0.04	0 (0.00)	—	0 (0.00)	8.71 (0.09)			9.63 (0.29)	0.20 (0.01)		
Truthful $L1$	0.14	0 (0.00)	—	0 (0.00)	8.71 (0.09)						
Truthful $L2$	0.01	0 (0.00)	—	0 (0.00)	8.71 (0.09)						
Equilibrium	0.16	0.02 (0.00)	29.84 (35.05)	0.01 (0.00)	8.71 (0.09)						
Log-likelihood	−569.53	−642.91		−644.12			−624.28	−684.81		−688.44	
BIC	−680.11	−655.92		−651.93			−728.31	−688.71		−689.74	

[a]This summary pools GHP's results for the low- and high-value treatments.

not nested, the BIC again favors subject-specific precisions. Given these results and that our primary purpose is to explore the models' ability to describe behavior rather than to forecast, we focus on the results for subject-specific precisions, with some attention to those for type-specific precisions and, in GHP, constant precisions.[36]

First consider the KL first-price estimates in Table IIIa. For the level-k plus equilibrium model with subject-specific precisions, we estimate 4% random $L0$, 59% random $L1$, 4% random $L2$, 18% truthful $L1$, and 16% truthful $L2$ or *Equilibrium* (not separated here) subjects. With type-specific precisions, we estimate 35% random $L1$, 3% random $L2$, 54% truthful $L1$, and 8% truthful $L2$ or *Equilibrium*.[37] The log-likelihood is substantially higher with subject-specific than type-specific or constant precisions, corresponding to the rejections via likelihood-ratio tests reported above. The BIC also favors the model with subject-specific precisions, but less strongly.

For the cursed-equilibrium models in the right half of Table IIIa, with subject-specific precisions and 11 cursed types, restricted to χ's that are multiples of 0.1 in $[0, 1]$, we estimate 47% of the subjects with $\chi = 1$ (fully-cursed equilibrium) and 20% with $\chi = 0$ (*Equilibrium*), with the remaining 33% spread almost uniformly over intervening values of χ. With type-specific precisions, the model estimates only three cursed types with positive frequency,

that the BIC does. For the common-value treatments, with type-specific precisions and random $L0$ plus four types (in each such treatment, one pair of types is not separated even by deviation costs), the level-k model has 8 independent parameters (4 type frequencies and 4 precisions). The analogous cursed-equilibrium model has 12 (4 levels of χ, 4 type frequencies, and 4 precisions; random $L0$ has no precision). For GHP, with type-specific precisions and random $L0$ plus five types, the level-k model has 10 independent parameters (5 type frequencies and 5 precisions). The analogous QRE model has 10 (5 type frequencies and 5 precisions). But as will be seen, for KL first-price and AK second-price, cursed-equilibrium models with type-specific and constant precisions estimate fewer types with positive frequencies than we allowed. For GHP, QRE models with type-specific and constant precisions also estimate fewer types than we allowed.

[36]Type-specific precisions, or a parameterized distribution of subject-specific precisions, are likely to be more useful for prediction. But models with subject-specific precisions are more robust to specification bias (e.g., if some subjects are very erratic but their precisions are constrained to equal those of other subjects) and so more useful as diagnostics. They also make our estimates more comparable with ER's, some of which allow subject-specific (though nonlogistic) error distributions. With subject-specific precisions, estimating (31) reduces to estimating subject by subject.

[37]In KL first-price with truthful types excluded from the specification, random $L1$ and $L2$ and in one case *Equilibrium* gain at the expense of truthful $L1$. With subject-specific precisions we then estimate 4% random $L0$ as before, 63% random $L1$, 18% random $L2$, and 16% *Equilibrium* as before. In this case, a restriction excluding truthful $L1$ and $L2$ is not rejected (p-value 0.2180). With type-specific precisions, we estimate 0% random $L0$ as before, 46% random $L1$, 6% random $L2$, and 48% *Equilibrium*; and a restriction excluding truthful $L1$ and $L2$ is strongly rejected (p-value 0.00098). With constant precisions, we estimate 0% random $L0$ as before, 53% random $L1$, 35% random $L2$, and 13% *Equilibrium*; and a restriction excluding truthful $L1$ and $L2$ is not rejected (p-value 0.2543).

$\chi = 0, 0.78$, and 0.99.[38] As for the level-k models, the log-likelihood is highest with subject-specific precisions, likelihood-ratio tests reject restrictions to type-specific or constant precisions (p-values far below 0.001), and the BIC favors subject-specific precisions. In this treatment (unlike AK or KL second-price), the constraint that χ is equal to either 0 or 1 is strongly rejected (p-value far below 0.001) and intermediate levels of χ fit some subjects better than random $L1$ ($\chi = 1$) or *Equilibrium* ($\chi = 0$).

Overall, in KL first-price, a cursed-equilibrium model has a modest likelihood advantage over a level-k model with subject-specific or type-specific precisions, which persists when the BIC is used to correct for its larger number of parameters with subject-specific precisions. Most (but not all) of the cursed-equilibrium model's advantage here is due to the fact that cursed types with intermediate values of χ fit some subjects better than any of our level-k types. (By contrast, in AK or KL second-price, intermediate values of χ add little to a cursed-equilibrium model's fit.)

Now consider the results for AK second-price in Table IIIc. For the level-k models, the estimated frequency of random $L0$ drops from the 4% estimated in KL first-price to 0% in AK second-price. As in most previous estimates from other settings, random $L0$ exists mainly in the minds of random $L1$ and (indirectly) of $L2$. With subject-specific precisions the estimated type frequencies are close to those for KL first-price. They are less close with type-specific precisions, but this may be due to restricting to type-specific precisions when they are strongly rejected.[39]

Turning to the cursed-equilibrium estimates in the right half of Table IIIc, with subject-specific precisions our cursed type frequency estimates are fairly close to those for KL first-price. But in AK second-price, unlike in KL first-price, our level-k models have substantial advantages in likelihood and the BIC over the cursed-equilibrium models for all error specifications.

Now consider the results for GHP first-price in Table IIId. For the level-k models, the estimated frequency of random $L0$ is again 0%. With subject-specific precisions, the estimated type frequencies are close to those for KL first-price and AK second-price. With type-specific precisions, the frequency

[38]When four cursed types are allowed, two of them are estimated to have $\chi = 1$, with different precisions; thus the extra type serves mainly to relax the restriction to type-specific precisions. Likelihood ratio tests fail to reject the restriction to three types (p-value 0.0639).

[39]In second-price common-value auctions, recalling that truthful $L1$ is equivalent to random $L2$ and truthful $L2$ is equivalent to random $L3$ (which we exclude for consistency), with truthful types excluded the estimates change mostly in interpretation. In AK second-price, random $L1$ and in one case *Equilibrium* gain at the expense of truthful $L2$. With subject-specific precisions, we estimate 0% random $L0$ as before, 74% random $L1$, 9% random $L2$, and 17% *Equilibrium*. A restriction excluding truthful $L2$ (with random $L2$ proxying for truthful $L1$) is strongly rejected (p-value 0.00045). With type-specific precisions, we estimate 0% random $L0$, 62% random $L1$, 0% random $L2$, and 38% *Equilibrium*. A restriction excluding truthful $L2$ is not rejected (p-value 0.2901). With constant precisions, the estimated frequency of truthful $L2$ was 0, so the estimates are unchanged.

estimates are less plausible (with a 98% frequency of random $L1$), which may again be due to imposing a restriction that is strongly rejected.[40] In GHP our level-k models have a substantial likelihood and BIC advantage over QRE models for all error specifications, and the model with subject-specific precisions gives a plausible account of some of the deviations from equilibrium that GHP attribute to risk aversion or joy of winning.

With subject-specific precisions, the estimated level-k type frequencies are very stable across KL first-price, AK second-price, and GHP: Ignoring lack of separation of types for the purpose of this comparison, the frequency of random $L1$ ranges from 0.59 to 0.65; of random $L2$ from 0.04 to 0.09; of truthful $L1$ from 0.09 to 0.18; of truthful $L2$ from 0.01 to 0.16; and of *equilibrium* from 0.04 to 0.16. (The estimated frequencies are less stable with type-specific precisions, but this may be due to imposing restrictions that are strongly rejected.) These type frequencies are behaviorally plausible and close to previous estimates (Stahl and Wilson (1995), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004), Costa-Gomes and Crawford (2006), Crawford and Iriberri (2007a)). However, the frequency of random $L1$ is higher than in most previous estimates, perhaps due to the heavier cognitive load of incomplete-information games.

The estimates for KL second-price in Table IIIb are very different. With subject-specific precisions, the estimated frequency of random $L1$, at 0.21, is far below the range of the other three treatments. The estimated frequency of truthful $L2$, at 0.32, is correspondingly high.[41] But in KL second-price as well, our level-k models have substantial advantages in likelihood and the BIC over cursed-equilibrium models for all error specifications.

We suggest a tentative explanation for the differences in the KL second-price estimates as follows.[42] In KL second-price, *equilibrium* shades its bid below the value suggested by its signal to adjust for the curse, random $L1$ bids the value

[40]With truthful types excluded in GHP first-price, random $L1$ and in one case *Equilibrium* gain at the expense of truthful $L1$ and $L2$. With subject-specific precisions, we estimate 0% random $L0$, 78% random $L1$, 4% random $L2$, and 19% *Equilibrium*; and a restriction excluding truthful $L1$ and $L2$ is rejected (p-value 0.0013). With type-specific or constant precisions, the estimated frequencies of truthful $L1$ and $L2$ were 0, so the estimates are unchanged.

[41]With truthful types excluded (and excluding random $L3$) in KL second-price, random $L1$ and in one case *Equilibrium* gain at the expense of truthful $L2$, just as in AK second-price. With subject-specific precisions, we estimate 18% random $L0$, 32% random $L1$, 21% random $L2$, and 29% *Equilibrium*. A restriction excluding truthful $L2$ is strongly rejected (p-value < 0.0000). With type-specific precisions, we estimate 49% random $L0$, 44% random $L1$, 10% random $L2$, and 0% *Equilibrium*. A restriction excluding truthful $L2$ is strongly rejected (p-value < 0.0000). We take the high estimated frequencies of random $L0$, unique to KL second-price with or without truthful types excluded, as probable evidence of misspecification. With constant precisions, we estimate 0% random $L0$, 0% random $L1$, 47% random $L2$, and 53% *Equilibrium*. A restriction excluding truthful $L2$ is strongly rejected (p-value < 0.0000).

[42]Similar deviations from the dominant bidding strategy occur in second-price independent-private-value auctions (Kagel, Harstad, and Levin (1987)). This and the evidence on experience and/or ability effects in Kagel and Richard (2001), Casari, Ham, and Kagel (2007), and Charness

suggested by its signal, truthful $L1$ and random $L2$ shade more than in equilibrium, and truthful $L2$ bids above the value suggested by its signal. There are two main patterns in the data: Some subjects shade their bids, but less than in equilibrium; in a level-k model, they are best captured by *Equilibrium* or random $L1$. Others bid above the values suggested by their signals; they are best captured by truthful $L2$. We suspect the latter subjects bid so high not because they believe (like truthful $L2$) that others are shading their bids more than in equilibrium, but because they do not fully process the subtle implications of the second-price auction for their optimal bidding strategy: They know they will not have to pay their own bid, and they underestimate its indirect cost via winning and paying more than the value, which may be less salient than what they will have to pay. Our model rules out this error by assumption, leaving truthful $L2$ as the best proxy for these subjects.

Turning in more detail to the cursed-equilibrium estimates in the right halves of Tables IIIa-c, despite our different specification and focus on inexperienced subjects, our cursed-equilibrium estimates for KL first-price and AK second-price are generally consistent with ER's estimates for their subjects, particularly AK's inexperienced subjects.[43] They are also close to our estimates for the level-k plus equilibrium model: For all three common-value treatments, with subject-specific precisions and 11 cursed types restricted to multiples of 0.1 in the interval [0, 1], there are spikes in the estimated distribution at $\chi = 1$ (fully-cursed equilibrium or random $L1$) and $\chi = 0$ (*Equilibrium*), and little weight on the intervening values (with minor exceptions at $\chi = 0.2$ in KL second-price and $\chi = 0.7$ in AK second-price). The estimated type frequencies for cursed types are similar except in KL second-price, where with type-specific precisions, the cursed-equilibrium model also breaks down, estimating the frequency of random $L0$ subjects as 0.43.

In all four treatments, the type frequency estimates have low standard errors, as do the precision estimates for the models with type-specific and constant precisions, with the exception of the precisions for types estimated to have very low population frequencies. However, the estimated precisions are highly heterogeneous across subjects; this is consistent with previous studies (e.g., Costa-Gomes and Crawford (2006)). Section D of the web appendix reports subject-specific precision estimates and their standard errors for our

and Levin (2005) suggest that some nonequilibrium bidding has nothing to do with strategic uncertainty and so cannot be explained by level-k thinking. This, and evidence from new experiments, should ultimately make it possible to build a more comprehensive model of bidders' behavior.

[43]In 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 for inexperienced subjects and to 0 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.

level-k plus equilibrium and cursed-equilibrium or QRE models in each treatment. To give a clearer picture of these estimates and how precisely they are estimated, Figures 1–8 present grouped histograms of the population distribution of precisions for each treatment–model pair: one of subjects' point precision estimates, one of subjects' estimates less their standard errors, and one of subjects' estimates plus their standard errors. The histograms show that the models with subject-specific precisions yield reasonably tight estimates of most subjects' precisions. To illustrate the implications of these and our type-specific and constant precision estimates for subjects' bidding behavior, Section E of the web appendix graphs the implied logit bid densities for each type with positive estimated frequency in each treatment, for representative low, medium, and high precision values.

To sum up, in KL and AK second-price, but not in KL first-price, our level-k models have substantial advantages in likelihood and the BIC over cursed-equilibrium models for all error specifications. In GHP, our level-k models also have a substantial likelihood and BIC advantage over QRE models, and they explain some of the deviations from equilibrium that GHP attribute to risk aversion or joy of winning. Thus, in two out of three common-value treatments, a level-k model with an empirically plausible type distribution fits the experimental data better than equilibrium or cursed equilibrium, and in GHP's independent-private-value treatments, a level-k model with a plausible type distribution fits the data better than equilibrium or QRE.

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 nonequilibrium 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, extend to an empirically plausible model of nonequilibrium bidding. The model yields tractable characterizations of the two factors that determine equilibrium bidding strategies in first- or second-price auctions: value adjustment for the information revealed by winning in common-value auctions (the winner's curse) and the bidding trade-off between the cost of higher bids and their higher probability of winning in first-price auctions with common or independent private values. These characterizations guide the choice of a model that can track the variation in subjects' initial responses to auctions across several experimental treatments.

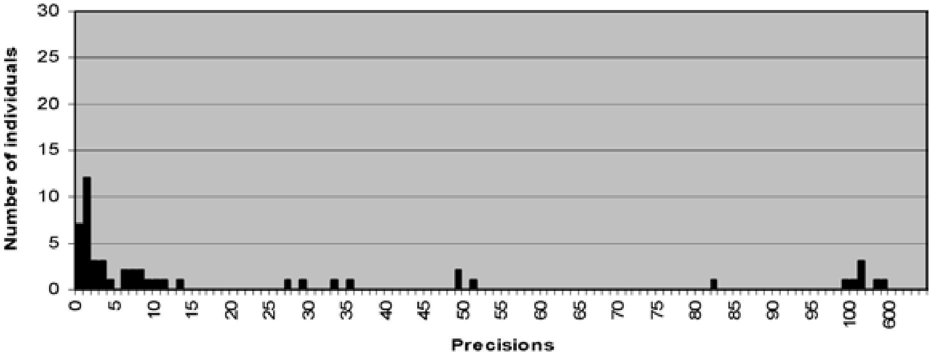

FIGURE 1A.—Subject-specific precisions in KL first-price level-k model.

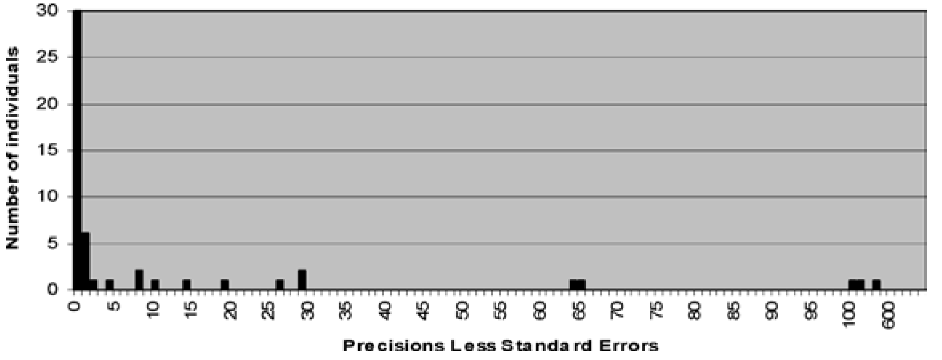

FIGURE 1B.—Subject-specific precisions less standard errors in KL first-price level-k model.

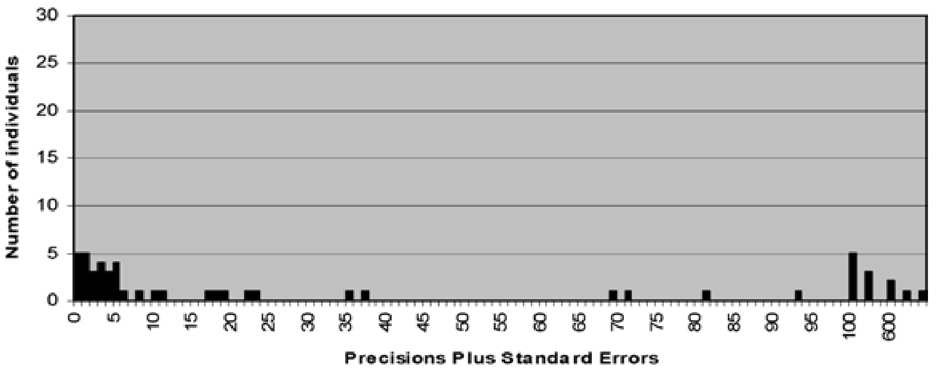

FIGURE 1C.—Subject-specific precisions plus standard errors in KL first-price level-k model.

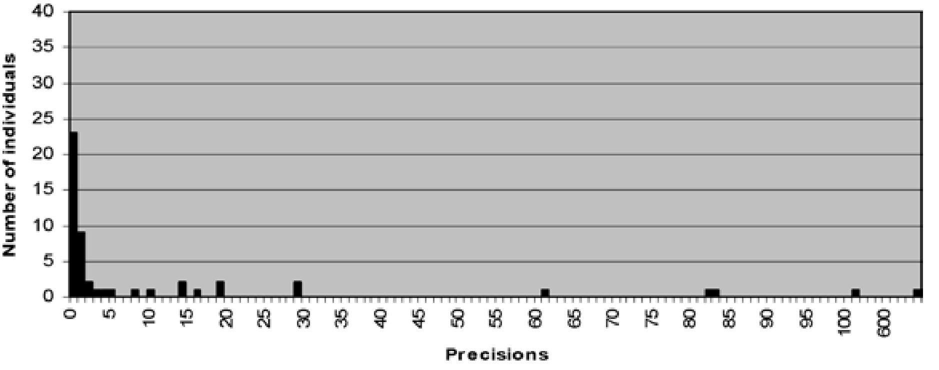

FIGURE 2A.—Subject-specific precisions in KL first-price cursed-equilibrium model.

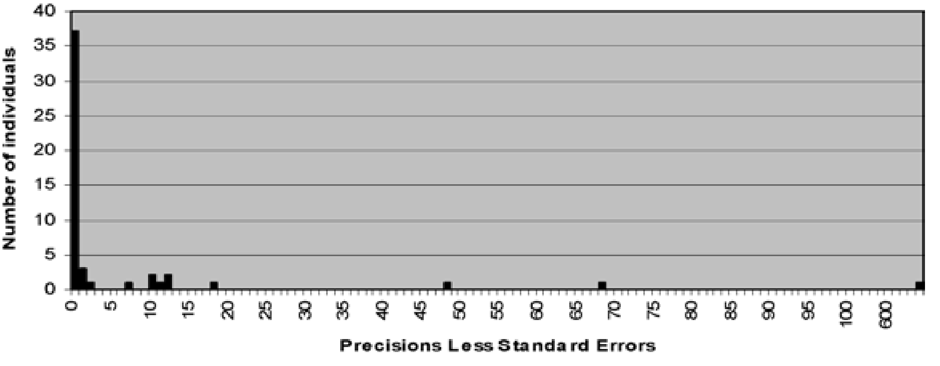

FIGURE 2B.—Subject-specific precisions less standard errors in KL first-price cursed-equilibrium model.

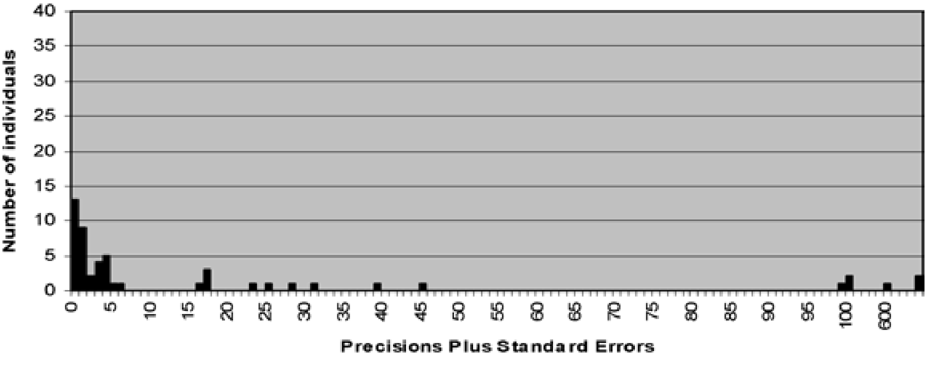

FIGURE 2C.—Subject-specific precisions plus standard errors in KL first-price cursed-equilibrium model.

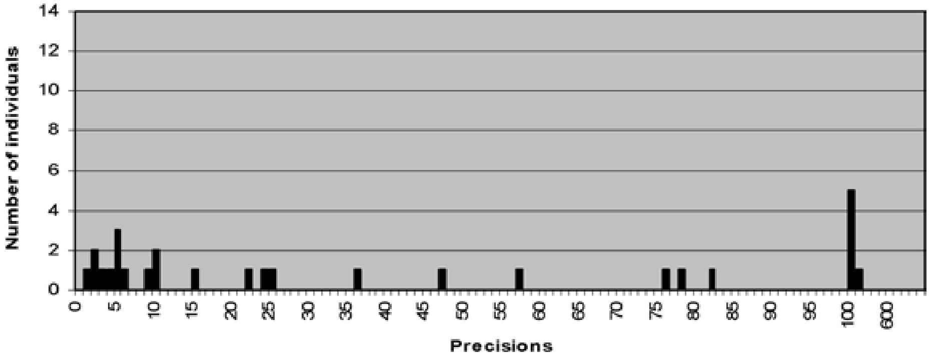

FIGURE 3A.—Subject-specific precisions in KL second-price level-*k* model.

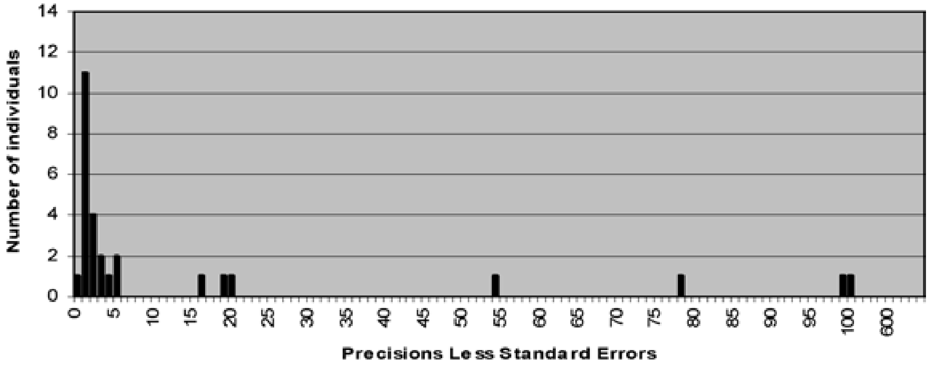

FIGURE 3B.—Subject-specific precisions less standard errors in KL second-price level-*k* model.

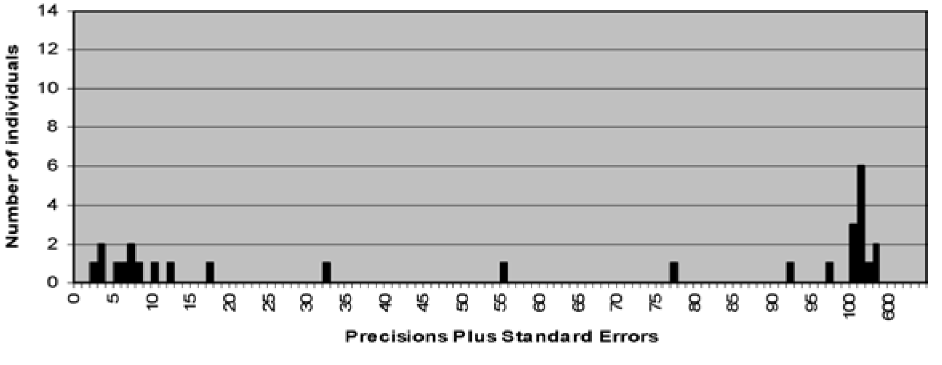

FIGURE 3C.—Subject-specific precisions plus standard errors in KL second-price level-*k* model.

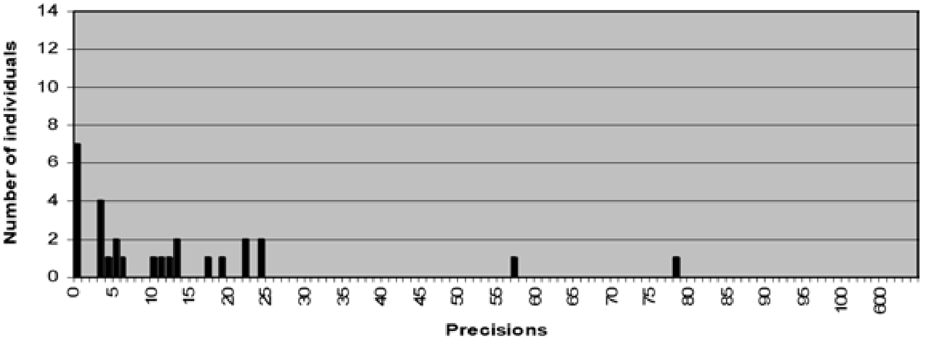

FIGURE 4A.—Subject-specific precisions in KL second-price cursed-equilibrium model.

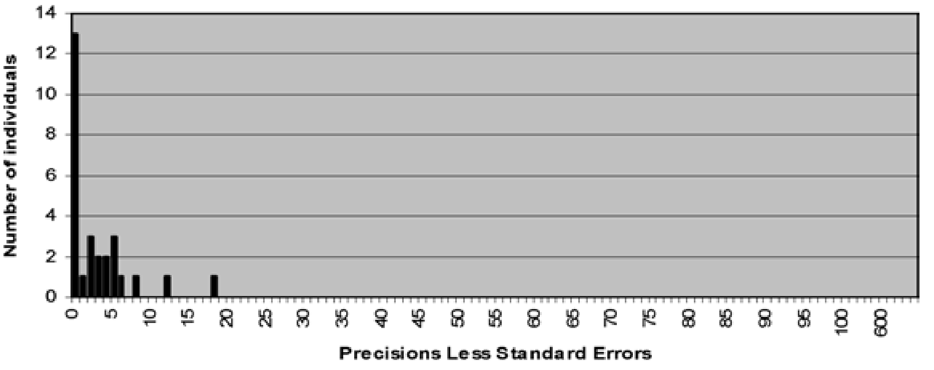

FIGURE 4B.—Subject-specific precisions less standard errors in KL second-price cursed-equilibrium model.

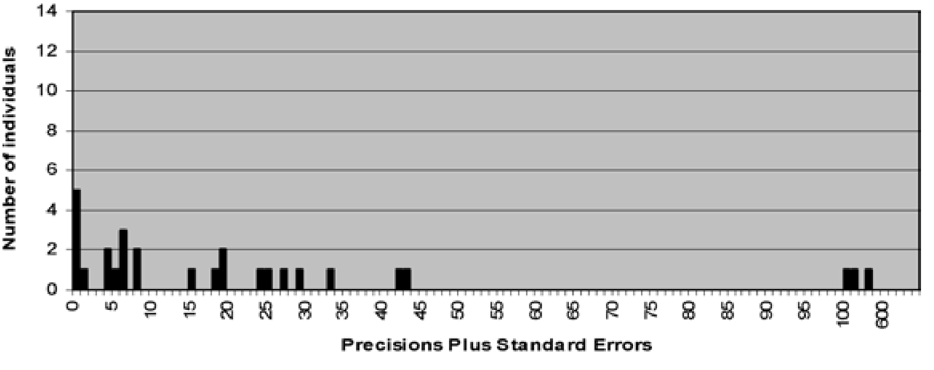

FIGURE 4C.—Subject-specific precisions plus standard errors in KL second-price cursed-equilibrium model.

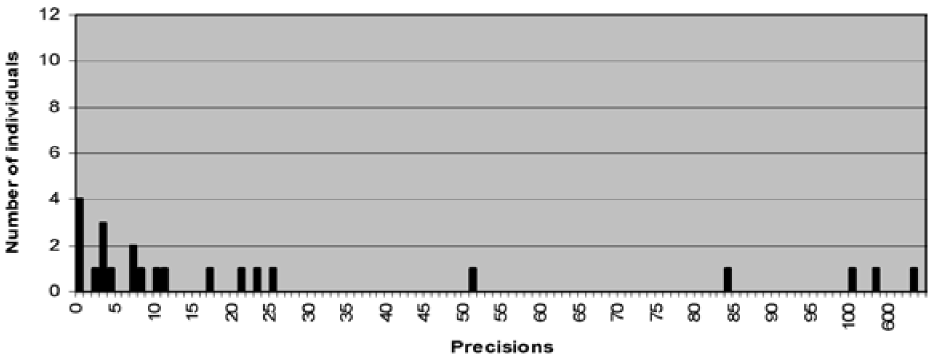

FIGURE 5A.—Subject-specific precisions in AK second-price level-*k* model.

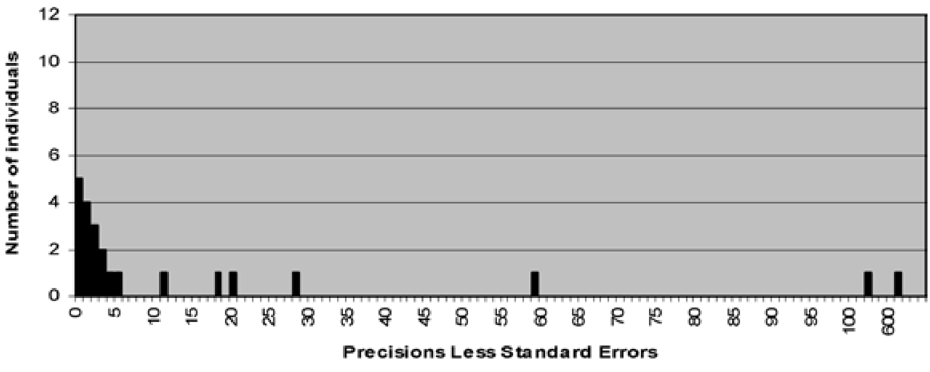

FIGURE 5B.—Subject-specific precisions less standard errors in AK second-price level-*k* model.

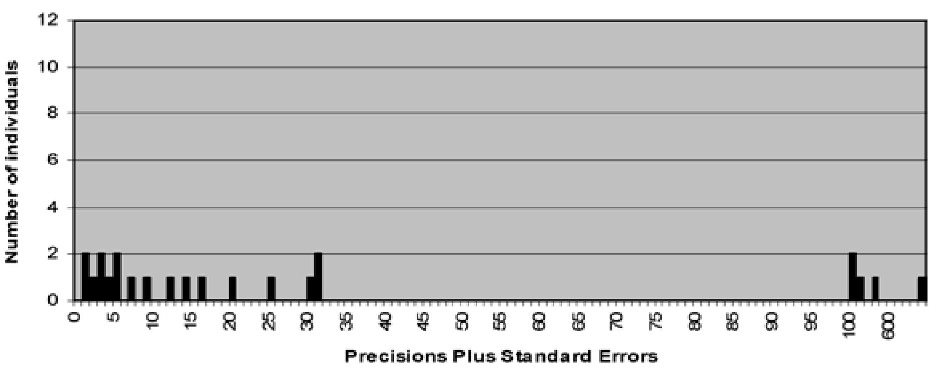

FIGURE 5C.—Subject-specific precisions plus standard errors in AK second-price level-*k* model.

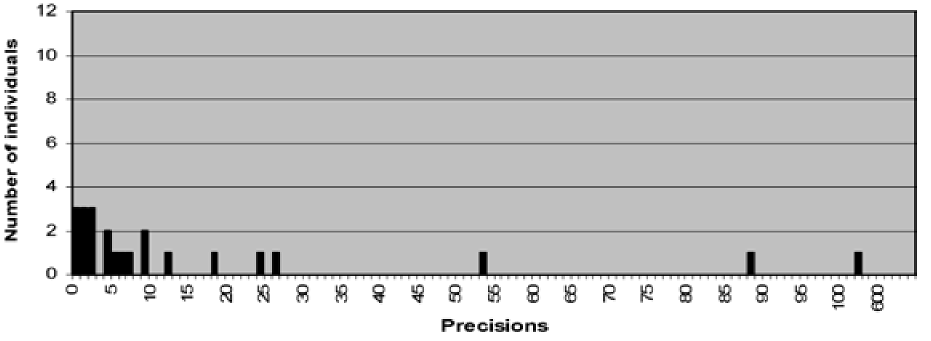

FIGURE 6A.—Subject-specific precisions in AK second-price cursed-equilibrium model.

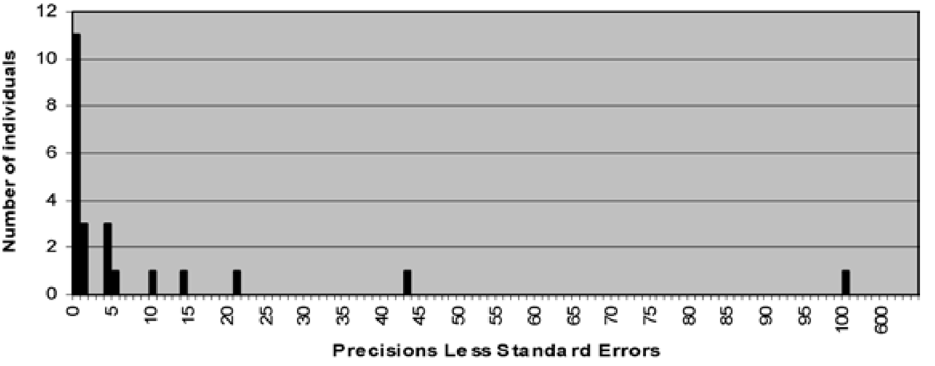

FIGURE 6B.—Subject-specific precisions less standard errors in AK second-price cursed-equilibrium model.

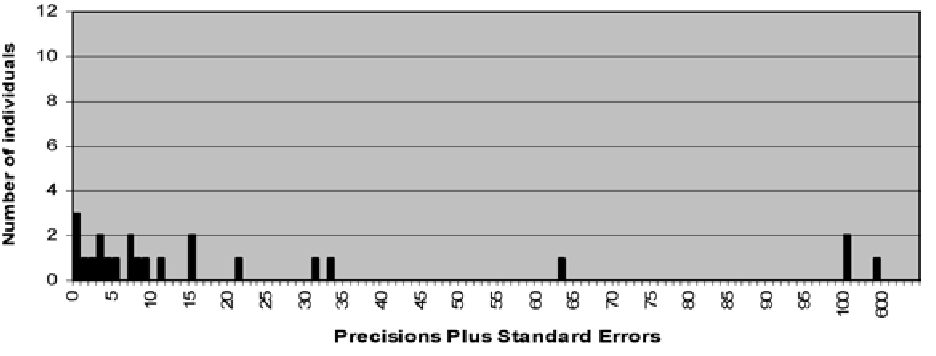

FIGURE 6C.—subject-specific precisions plus standard errors in AK second-price cursed-equilibrium model.

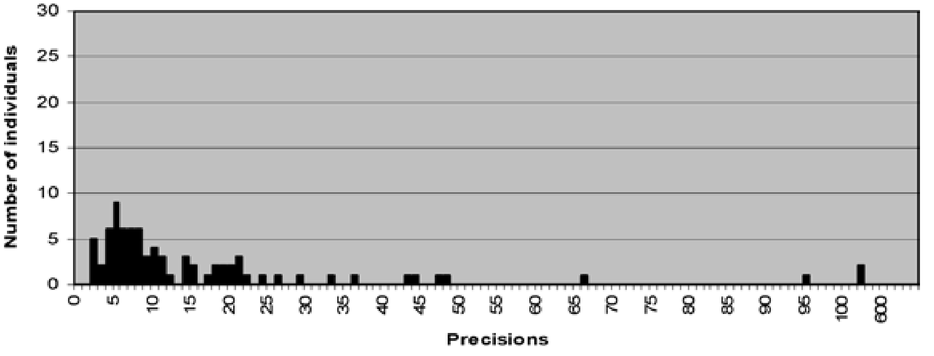

FIGURE 7A.—Subject-specific precisions in GHP first-price level-k model.

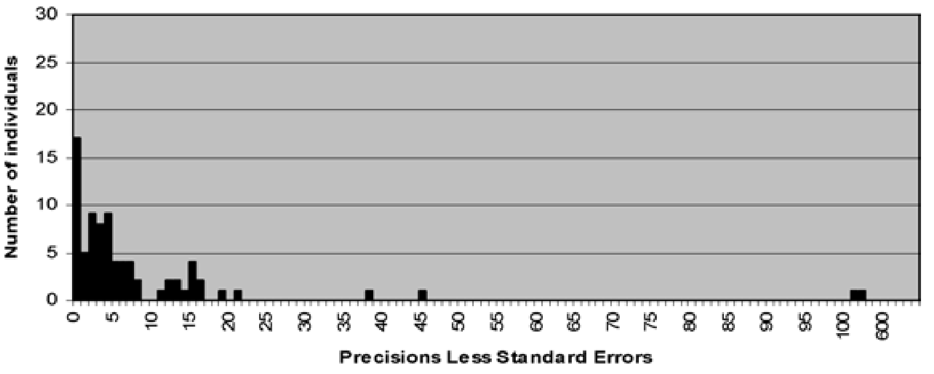

FIGURE 7B.—Subject-specific precisions less standard errors in GHP first-price level-k model.

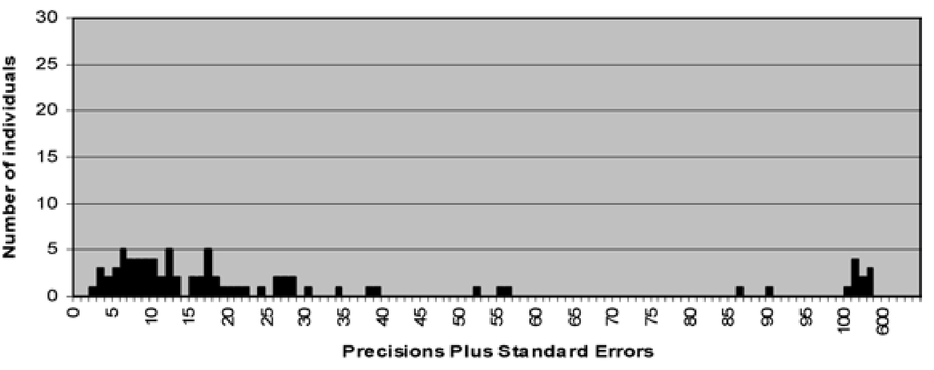

FIGURE 7C.—Subject-specific precisions plus standard errors in GHP first-price level-k model.

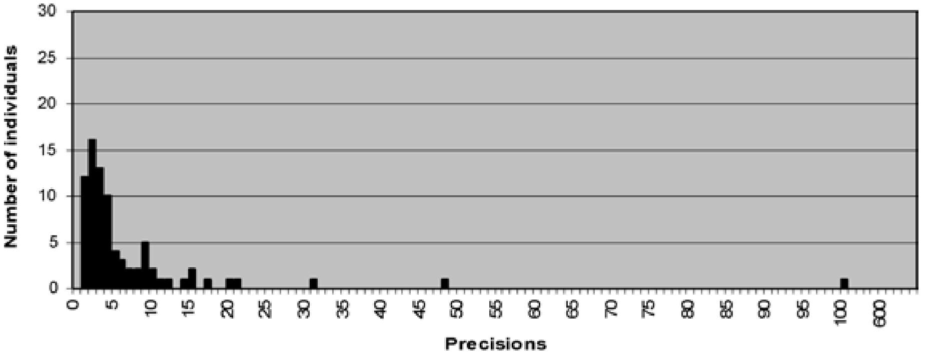

FIGURE 8A.—Subject-specific precisions in GHP QRE first-price model.

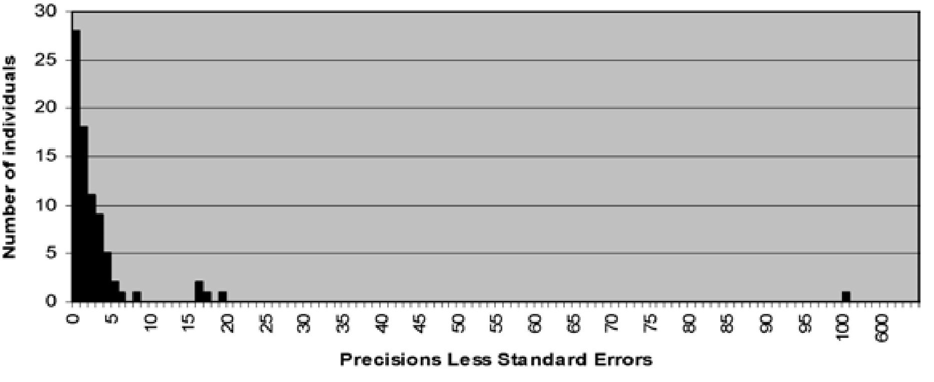

FIGURE 8B.—Subject-specific precisions less standard errors in GHP first-price QRE model.

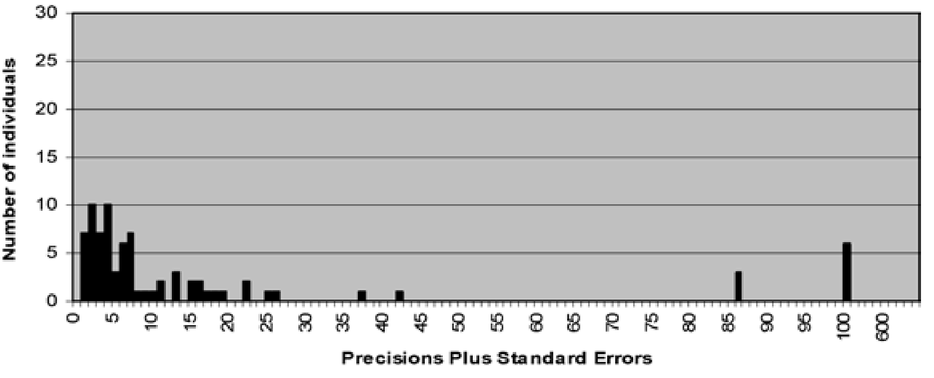

FIGURE 8C.—Subject-specific precisions plus standard errors in GHP first-price QRE model.

In our econometric analysis, a level-k model with an empirically plausible type distribution fits better except in KL first-price than the leading alternatives of cursed equilibrium or QRE, and yields a simple, unified explanation of the winner's curse in some leading common-value auction designs and overbidding in some independent-private-value auction designs with nonuniform value distributions. Random $L1$ is by far the most frequent type in all but KL second-price, with truthful $L1$ playing a substantial supporting role. Thus most subjects' behavior is strategic, even though it does not usually conform to equilibrium. Even though random $L1$ yields the same bidding strategies as Eyster and Rabin's notion of fully-cursed equilibrium in the common-value treatments, our estimated level-k type distribution fits the distribution of subjects' responses better than an estimated model with the same number of cursed types in all but KL first-price.

Thus, by viewing behavior in these auctions through the lens of a general, portable model of strategic behavior, the level-k model allows us to link a large body of data from auction experiments, most of which has been analyzed by assuming equilibrium in some form, to data from nonauction experiments that were specifically designed to study strategic thinking.

Dept. of Economics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0508, U.S.A.; vcrawfor@dss.ucsd.edu

and

Dept. of Economics and Business, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain; nagore.iriberri@upf.edu.

Manuscript received October, 2005; final revision received June, 2007.

REFERENCES

AVERY, C., AND J. H. KAGEL (1997): "Second-Price Auctions with Asymmetric Payoffs: An Experimental Investigation," *Journal of Economics and Management Strategy*, 6, 573–603. [1721,1728,1729,1740,1746]

BALL, S., M. H. BAZERMAN, AND J. CARROLL (1991): "An Evaluation of Learning in the Bilateral Winner's Curse," *Organizational Behavior and Human Decision Processes*, 48, 1–22. [1724]

BATTIGALLI, P., AND M. SINISCALCHI (2003): "Rationalizable Bidding in First-Price Auctions," *Games and Economic Behavior*, 45, 38–72. [1724]

BAZERMAN, M. H., AND W. F. SAMUELSON (1983): "I Won the Auction But I Don't Want the Prize," *Journal of Conflict Resolution*, 27, 618–634. [1722]

BOSCH-DOMÈNECH, A., J. G. MONTALVO, R. NAGEL, AND A. SATORRA (2002): "One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, 92, 1687–1701. [1724]

BRANDTS, J., AND C. HOLT (1992): "An Experimental Test of Equilibrium Dominance in Signaling Games," *American Economic Review*, 82, 1350–1365. [1734]

——— (1993): "Adjustment Patterns and Equilibrium Selection in Experimental Signaling Games," *International Journal of Game Theory*, 22, 279–302. [1734]

CAI, H., AND J. T. WANG (2006): "Overcommunication in Strategic Information Transmission Games," *Games and Economic Behavior*, 56, 7–36. [1734]

CAMERER, C., T.-H. HO, AND J.-K. CHONG (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119, 861–898. [1724-1726,1734,1746,1756]

CAPEN, E. C., R. V. CLAPP, AND W. M. CAMPBELL (1971): "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 23, 641–653. [1722]
CASARI, M., J. HAM, AND J. KAGEL (2007): "Selection Bias, Demographic Effects and Ability Effects in Common Value Auction Experiments," *American Economic Review*, 97, 1278–1304. [1756]
CHARNESS, G., AND D. LEVIN (2005): "The Origin of the Winner's Curse: A Laboratory Study," Unpublished Manuscript, University of California, Santa Barbara. [1722,1725,1756,1757]
CHEN, K.-Y., AND C. R. PLOTT (1998): "Nonlinear Behavior in Sealed Bid First Price Auctions," *Games and Economic Behavior*, 25, 34–78. [1746]
COMPTE, O. (2004): "Prediction Errors and the Winner's Curse," Unpublished Manuscript, CERAS, École Nationale des Ponts et Chausses. [1726]
COSTA-GOMES, M. A., AND V. P. CRAWFORD (2006): "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, 96, 1737–1768. [1724-1726,1733,1746,1756,1757]
COSTA-GOMES, M. A., V. P. CRAWFORD, AND B. BROSETA (2001): "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 69, 1193–1235. [1724,1726, 1746,1756]
COX, J. C., V. L. SMITH, AND J. M. WALKER (1983): "Tests of a Heterogeneous Bidders Theory of First Price Auctions," *Economics Letters*, 12, 207–212. [1722,1723]
——— (1988): "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, 1, 61–99. [1722,1723]
——— (1992): "Theory and Misbehavior of First-Price Auctions: Comment," *American Economic Review*, 82, 1392–1412. [1723]
CRAWFORD, V. P. (1998): "A Survey of Experiments on Communication via Cheap Talk," *Journal of Economic Theory*, 78, 286–298. [1734]
——— (2003): "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review*, 93, 133–149. [1724,1734]
CRAWFORD, V. P., AND N. IRIBERRI (2007a): "Fatal Attraction: Salience, Naivete, and Sophistication in Experimental 'Hide-and-Seek' Games," *American Economic Review*, 97, forthcoming. [1724,1734,1746,1756]
——— (2007b): "Supplement to 'Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-value Auctions?'," *Econometrica Supplementary Material*, 75, http://www.econometricsociety.org/ecta/Supmat/6126_extensions.pdf; http://www.econometricsociety.org/ecta/Supmat/6126_data and programs.zip; http://www.econometricsociety.org/ecta/Supmat/6126_tables-figures.pdf. [1737, 1749]
DORSEY, R., AND L. RAZZOLINI (2003): "Explaining Overbidding in First Price Auctions Using Controlled Lotteries," *Experimental Economics*, 6, 123–140. [1725]
ESPONDA, I. (2005): "Behavioral Equilibrium in Economies with Adverse Selection," Unpublished Manuscript, Stanford University. [1723]
EYSTER, E., AND M. RABIN (2002): "Cursed Equilibrium," Working Paper E02-320, IBER, UC Berkeley. [1722,1728,1731]
——— (2005): "Cursed Equilibrium," *Econometrica*, 73, 1623–1672. [1722,1728,1731,1732, 1741,1747,1758]
FEDDERSEN, T., AND W. PESENDORFER (1996): "The Swing Voter's Curse," *American Economic Review*, 86, 408–424. [1722]
——— (1997): "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica*, 65, 1029–1058. [1722]
——— (1998): "Convicting the Innocent: the Inferiority of Unanimous Jury Verdicts under Strategic Voting," *American Political Science Review*, 90, 23–35. [1722]
FILIZ-OZBAY, E., AND E. Y. OZBAY (2007): "Auctions with Anticipated Regret: Theory and Experiment," *American Economic Review*, 97, 1407–1418. [1724]
FRIEDMAN, D. (1998): "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly," *American Economic Review*, 88, 933–946. [1722]

FUDENBERG, D. (2006): "Advancing Beyond 'Advances in Behavioral Economics'," *Journal of Economic Literature*, 44, 694–711. [1724]

FUDENBERG, D., AND D. K. LEVINE (1993): "Self-Confirming Equilibrium," *Econometrica*, 61, 523–545. [1723]

GARVIN, S., AND J. H. KAGEL (1994): "Learning in Common Value Auctions: Some Initial Observations," *Journal of Economic Behavior and Organization*, 25, 351–372. [1721,1724,1728,1746]

GNEEZY, U. (2005): "Step-Level Reasoning and Bidding in Auctions," *Management Science*, 51, 1633–1642. [1726]

GOEREE, J. K., AND C. A. HOLT (2001): "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91, 1402–1422. [1726,1746]

GOEREE, J. K., C. A. HOLT, AND T. R. PALFREY (2002): "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, 104, 247–272. [1722,1723, 1728,1741,1743,1744,1746]

HENDRICKS, K., AND R. H. PORTER (1988): "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 78, 865–883. [1722]

HENDRICKS, K., R. H. PORTER, AND B. BOUDREAU (1987): "Information, Returns, and Bidding Behavior in OCS Auctions: 1954–1969," *Journal of Industrial Economics*, 35, 517–542. [1722]

HO, T.-H., C. CAMERER, AND K. WEIGELT (1998): "Iterated Dominance and Iterated Best Response in Experimental p-Beauty Contests," *American Economic Review*, 88, 947–969. [1724,1734]

HOLT, C. A., AND R. SHERMAN (1994): "The Loser's Curse," *American Economic Review*, 84, 642–652. [1722,1723]

——— (2000): "Risk Aversion and the Winner's Curse," Working Paper, University of Virginia. [1723,1746]

JEHIEL, P. (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, 123, 81–104. [1723]

JEHIEL, P., AND F. KOESSLER (2005): "Revisiting Games of Incomplete Information with Analogy-Based Expectations," Unpublished Manuscript, CERAS, École Nationale des Ponts et Chausses. [1723]

KAGEL, J. H. (1995): "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel and A. E. Roth. Princeton, NJ: Princeton University Press, 501–585. [1722]

KAGEL, J. H., AND D. LEVIN (1986): "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76, 894–920. [1721,1722,1728,1729,1740,1746]

——— (1993): "Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders," *Economic Journal*, 103, 868–879. [1746]

——— (2002): *Common Value Auctions and the Winner's Curse*. Princeton, NJ: Princeton University Press, [1721,1722,1746]

KAGEL, J. H., AND J.-F. 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. [1724,1756]

KAGEL, J. H., R. HARSTAD, AND D. LEVIN (1987): "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study," *Econometrica*, 55, 1275–1304. [1722,1756]

KAGEL, J. H., D. LEVIN, AND R. HARSTAD (1995): "Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-Price Common Value Auctions," *International Journal of Game Theory*, 24, 293–319. [1721,1722,1746]

KLEMPERER, P. (1999): "Auction Theory: A Guide to the Literature," *Journal of Economic Surveys*, 13, 277–286. [1721]

LAFFONT, J.-J. (1997): "Game Theory and Empirical Economics: The Case of Auction Data," *European Economic Review*, 41, 1–35. [1722]

LIND, B., AND C. R. PLOTT (1991): "The Winner's Curse: Experiments with Buyers and with Sellers," *American Economic Review*, 81, 335–346. [1722-1724,1729,1734,1740,1746]

MASKIN, E., AND J. RILEY (2000): "Asymmetric Auctions," *Review of Economic Studies*, 67, 413–438. [1726]

MCAFEE, R. P., AND J. MCMILLAN (1987): "Auctions and Bidding," *Journal of Economic Literature*, 25, 699–738. [1721,1722,1726]

MCKELVEY, R., AND T. PALFREY (1995): "Quantal Response Equilibria for Normal-Form Games," *Games and Economic Behavior*, 10, 6–38. [1723]

MILGROM, P. (1985): "The Economics of Competitive Bidding: A Selective Survey," in *Social Goals and Social Organization: A Volume in Honor of Elisha Pazner*, ed. by L. Hurwicz, D. Schmeidler, and H. Sonnenschein. London: Cambridge University Press, 261–289. [1721]

——— (1987): "Auction Theory," in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley. Cambridge, U.K.: Cambridge University Press, 1–32. [1721]

MILGROM, P., AND R. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122. [1721,1728-1731,1737]

NAGEL, R. (1995): "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313–1326. [1724]

NEUGEBAUER, T., AND R. SELTEN (2006): "Individual Behavior of First-Price Auctions: The Importance of Information Feedback in Computerized Experimental Markets," *Games and Economic Behavior*, 54, 183–204. [1724]

PALACIOS-HUERTA, I. (2003): "Learning to Open Monty Hall's Doors," *Experimental Economics*, 6, 235–251. [1722,1724]

PALFREY, T. (1985): "Buyer Behavior and the Welfare Effects of Bundling by a Multiproduct Monopolist: A Laboratory Test," in *Research in Experimental Economics*, Vol. 3. ed. by V. Smith. Greenwich, CT: JAI Press, 73–104. [1746]

PARTOW, J., AND A. SCHOTTER (1993): "Does Game Theory Predict Well for the Wrong Reasons? An Experimental Investigation," Working Paper 93-46, C. V. Starr Center for Applied Economics, New York University. [1734]

SAMUELSON, W. F., AND M. H. BAZERMAN (1985): "The Winner's Curse in Bilateral Negotiations," in *Research in Experimental Economics*, Vol. 3, ed. by V. L. Smith. Greenwich, CT: JAI Press, 105–137. [1722,1723]

SELTEN, R. (1998): "Features of Experimentally Observed Bounded Rationality," *European Economic Review*, 42, 413–436. [1725]

SONSINO, D., I. EREV, AND S. GILAT (2002): "On the Likelihood of Repeated Zero-Sum Betting by Adaptive (Human) Agents," Unpublished Manuscript, Technion, Israel Institute of Technology. [1722,1734]

SOVIK, Y. (2000): "Impossible Bets: An Experimental Study," Unpublished Manuscript, University of Oslo. [1722,1734]

STAHL, D., AND P. WILSON (1994): "Experimental Evidence on Players' Models of Other Players," *Journal of Economic Behavior and Organization*, 25, 309–327. [1724-1726,1746]

——— (1995): "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 218–254. [1724-1726,1746,1756]

THALER, R. H. (1988): "The Winner's Curse," *Journal of Economic Perspectives*, 2, 191–202. [1722]

TOR, A., AND M. H. BAZERMAN (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. [1722]

WEIZSÄCKER, G. (2003): "Ignoring the Rationality of Others: Evidence from Experimental Normal-Form Games," *Games and Economic Behavior*, 44, 145–171. [1731]

WILSON, R. (1969): "Competitive Bidding with Disparate Information," *Management Science*, 15, 446–448. [1722]

——— (1992): "Strategic Analysis of Auctions," in *Handbook of Game Theory*, Vol. 1, ed. by R. Aumann and S. Hart. Amsterdam: North-Holland, 227–279. [1721,1722]