Econometrica, Vol. 73, No. 5 (September, 2005), 1623–1672

CURSED EQUILIBRIUM

BY ERIK EYSTER AND MATTHEW RABIN[1]

There is evidence that people do not fully take into account how other people's actions depend on these other people's information. This paper defines and applies a new equilibrium concept in games with private information, *cursed equilibrium*, which assumes that each player correctly predicts the distribution of other players' actions, but underestimates the degree to which these actions are correlated with other players' information. We apply the concept to common-values auctions, where cursed equilibrium captures the widely observed phenomenon of the winner's curse, and to bilateral trade, where cursedness predicts trade in adverse-selections settings for which conventional analysis predicts no trade. We also apply cursed equilibrium to voting and signalling models. We test a single-parameter variant of our model that embeds Bayesian Nash equilibrium as a special case and find that parameter values that correspond to cursedness fit a broad range of experimental datasets better than the parameter value that corresponds to Bayesian Nash equilibrium.

KEYWORDS: Adverse selection, winner's curse, common-values auctions, speculative trade.

1. INTRODUCTION

A WIDELY OBSERVED PHENOMENON in laboratory auctions is the "winner's curse": when bidders who share a common but unknown value for a good have private information about the good's value, they tend to bid more than equilibrium theory predicts. In many experiments, the average winning bid exceeds the average value of the good. One explanation is that the typical bidder fails to fully appreciate that the low bids by other bidders needed for her to win the auction mean that these other bidders have private information that is more negative than her own. This failure leads the bidder to believe that the value of the object when she wins the auction is closer to the value suggested by her private information than it actually is and, hence, to overbid. Fully rational bidders avoid this problem by tempering their bids.

While the winner's curse has been observed repeatedly in laboratory experiments, and anecdotes and some research suggests that it is important outside of the laboratory, theoretical research on auctions assumes that people do not

[1]We thank Colin Camerer and seminar participants at Arizona, Berkeley, Birmingham, Caltech, Edinburgh, ESSET 2001 (Gerzensee), ESA Meeting 2002 (Strasbourg), Essex, Harvard, LSE, Nottingham, Nuffield, UCL, Zurich, Pompeu Fabra, and especially Dirk Engelmann, three anonymous referees, and Glenn Ellison for helpful comments, as well as Davis Beekman, Kitt Carpenter, David Huffman, and extra especially Jeff Holman for valuable research assistance. We are grateful to Chris Avery, John Forsythe, Serena Guarnaschelli, John Kagel, Mark Isaac, Dan Levin, Richard McKelvey, and Tom Palfrey for sharing their data. Eyster thanks the Olin and MacArthur Foundations, and Rabin thanks the Russell Sage, MacArthur, and National Science Foundations for financial support.

make this error.[2] Indeed, to estimate bidders' valuations, most empirical research explicitly assumes that bidders do not succumb to the winner's curse. Kagel and Levin (1986) and others in the context of common-values auctions, as well as Holt and Sherman (1994) in the context of trade with adverse selection, have posited and tested an extreme form of the winner's curse: agents act as if there is no information content in winning an auction or completing a trade.[3] In this paper, we formally model a generalization of the winner's curse that assumes that players in a Bayesian game underestimate the extent to which other players' actions are correlated with their information. Our model generalizes those of Kagel and Levin (1986) and Holt and Sherman (1994) both by allowing players to *partially* appreciate the information content in other players' actions and through its applicability to general Bayesian games. We flesh out its implications in common-values auctions and many other settings, arguing that it ties a wide range of empirically observed phenomena together with a formalization of a single psychological principle—the underappreciation of the informational content of other people's behavior. We illustrate the model's empirical fit with several experimental datasets and discuss how it captures aspects of behavior not captured by prominent existing generalizations of Bayesian Nash equilibrium. While we develop variants of the model that are more realistic and better accommodate the data, we primarily emphasize a simple single-parameter model that embeds Bayesian equilibrium as a particular parameter value and suggest that a single value of this parameter that corresponds to cursedness is more likely to fit a broad range of data than the value that corresponds to Bayesian Nash equilibrium.

In Section 2 we present our equilibrium concept. We consider standard Bayesian games where players' private information is represented by their payoff types, with the restriction that the set of strategies available to any given player does not depend on her type. Our equilibrium concept, *cursed equilibrium*, assumes that each player incorrectly believes that with positive probability each profile of types of the other players plays the same mixed action profile that corresponds to their average distribution of actions, rather than their true, type-specific action profile. Players choose their actions to maximize their expected payoffs given their types and these incorrect beliefs about other players' equilibrium strategies. We parameterize the extent to which a player is "cursed" by the probability $\chi \in [0, 1]$ she assigns to other players playing their average distribution of actions irrespective of type rather than their type-contingent strategy, to which she assigns probability $1 - \chi$. Setting $\chi = 0$ corresponds to the fully rational Bayesian Nash equilibrium and setting $\chi = 1$ corresponds to the case where each player assumes no connection whatsoever

[2] See Thaler (1988) for an overview of the early evidence on the winner's curse, as well as Kagel (1995) and Kagel and Levin (2002) for surveys of laboratory auctions.

[3] Potters and Wit (1995) and Jacobsen, Potters, Schram, van Winden, and Wit (2000) use this same premise to analyze markets for assets whose values are common but unknown to the traders.

between other players' actions and their types. Whatever χ, each player correctly predicts the equilibrium distribution of the other players' actions—the players' only mistake is underappreciating the connection between other players' types and their actions.

To illustrate cursed equilibrium, consider a simple variant of Akerlof's (1970) lemons model in which a buyer might purchase a car from a seller at a predetermined price of \$1,000 (Canadian). The seller knows whether the car is a lemon, worth \$0 (Canadian) to both seller and buyer, or a peach, worth \$3,000 (Canadian) to the buyer and \$2,000 (Canadian) to the seller. The buyer believes each occurs with probability $\frac{1}{2}$. The parties simultaneously announce whether they wish to trade, and the car is sold if and only if both say they wish to trade. While a fully rational buyer would realize that the seller wishes to trade if and only if the car is a lemon, and hence refuse to buy, a cursed buyer may mistakenly buy the car. A χ-cursed buyer believes that with probability χ the seller sells with probability $\frac{1}{2}$ irrespective of the type of car, so that the car being sold is a peach with probability $(1-\chi)\cdot 0+\chi\cdot\frac{1}{2}=\frac{\chi}{2}$, and therefore worth $\frac{\chi}{2}\cdot 3{,}000=1{,}500\chi$. Hence, a buyer cursed by $\chi>\frac{2}{3}$ will wish to buy the car, only to discover that whenever the wish comes true, the car is a worthless lemon.[4]

We prove that every finite game has (for every value of χ) a cursed equilibrium. We also show that when no player's payoffs depend on any other players' types, cursed equilibrium and Bayesian Nash equilibrium coincide. Intuitively, the only difference between the two concepts is that in a cursed equilibrium players have incorrect beliefs about the relationship between their opponents' actions and their types; if no player's payoffs depend on any other player's type, such mistaken beliefs do not matter. At the end of Section 2, we discuss various motivations for cursed equilibrium in more depth and compare it to other models of departures from Bayesian Nash equilibrium.

In Sections 3 and 4 we apply the general model to two different economic settings—bilateral trade and auctions. Our model both helps explain existing experimental behavior in these settings and provides plausible, testable predictions in settings for which we know of no experimental evidence. In Section 3 we examine adverse selection and no-trade theorems in the context of bilateral trade. When, as in the example above, a seller has private information about the value of a good, while the buyer does not, cursed equilibrium may lead to more trade than Bayesian Nash equilibrium: when only sellers with low-value goods sell, a buyer who fails to recognize this may buy when she would be better off not buying. Cursed equilibrium may also lead to less trade than Bayesian Nash equilibrium: because a cursed buyer does not fully appreciate that sellers with high-value goods sell at high prices, she may be too reluctant to

[4]This result suggests that a more general implication of our model is that people should be careful what they wish for. See also von Goethe (1808, 1832; 1949, 1959) and Wilde (1891, 1982), who obtained similar results in different frameworks.

pay higher prices. We show that the predictions of cursed equilibrium approximately correspond to the behavior of subjects in experimental tests of a lemons model by Samuelson and Bazerman (1985) and Holt and Sherman (1994). We also illustrate how in a setting with two-sided private information and common preferences, both parties may strictly prefer trading to not trading, in contrast to "no-trade results" such as those presented in Milgrom and Stokey (1982). This is because a buyer or seller who underinfers the other party's information conditional on trade may agree to a trade with a negative expected value. While trade may occur, it cannot occur with probability 1, which distinguishes our theory from others including noncommon priors.

In Section 4 we turn to our primary motivating application, common-values auctions. In a cursed equilibrium, bidders do not fully appreciate the information about the object's value conveyed by winning. Depending on their signals about the object's value, this may lead bidders to either increase or decrease their bids. Usually—but not always—the seller's expected revenue increases in χ. We show that when the number of bidders is large, bidders suffer the winner's curse—the average winning bid exceeds the average value of the object. Conversely, given a symmetry assumption, we establish that cursedness can only give rise to the winner's curse in auctions with four or more bidders. Finally, we compare the predictions of cursed equilibrium to some of the experimental evidence on common-values auctions.

In Section 5 we more briefly study the implications of cursed equilibrium to a range of applications, emphasizing how cursedness affects information revelation between parties. We contrast our model's predictions to those of a recent rational-choice literature on voting in elections and on juries that assumes people vote based on the information content in being pivotal. Because they underinfer the information content in being pivotal, cursed voters more naively vote according to their beliefs at the time of voting. This, in turn, implies that in contrast to the rational-choice literature, voting rules in large elections matter in a cursed equilibrium: whereas uncursed voters adjust their behavior to the voting rule to assure the efficient outcome, sufficiently cursed voters do not react to voting rules, so that rules are efficient if and only if they implement the right outcome when voters vote naively. Analysis of data from an experiment by Guarnaschelli, McKelvey, and Palfrey (2000) provides no direct support for our model with small juries (but support with large juries), but we are able to show that controlling for the cost of their errors, subjects make errors consistent with cursedness more frequently than those inconsistent with cursedness.

We then illustrate the implications of cursed equilibrium in two additional signalling contexts. First, we consider classical simple signalling games, where fully cursed equilibrium rules out the use of costly signalling, but lesser degrees of cursedness can either destroy or *create* meaningful signalling. Second, we apply cursed equilibrium to a model of "verifiable cheap talk" modelled after American political elections where voters make inferences about candidates after these candidates strategically reveal or conceal information about their past indiscretions or future plans. In this game, the unique perfect Bayesian Nash equilibrium calls for each type of politician to reveal her type, since fully rational voters would infer the worst from silence. Because cursed voters may not infer the worst from silence—they may believe that even "good" types conceal—even politicians with not-so-bad information may not reveal the truth. Data from an experiment by Forsythe, Isaac, and Palfrey (1989) support this prediction.

We believe that cursed equilibrium provides a simple, tractable, and portable psychologically motivated alternative to Bayesian Nash equilibrium that can improve empirical analysis in a broad array of settings. Indeed, while surely a more thorough search would produce counterexamples, with the exception of one condition in one experiment, any value of $\chi \in (0, 0.6)$ provides a better fit than does Bayesian Nash equilibrium in all the experiments we analyze. Nonetheless, shortcomings of the model limit its applicability beyond the set of games we consider. In some contexts, it makes some unrealistic predictions; in others, it is simply not well defined. Its implications can also depend on (seemingly artificial) ways to reformulate a Bayesian game, making it less robust than ideal for a solution concept. We conclude in Section 6 with a discussion of possible extensions that address these problems as well as some further economic applications of the principles developed in the paper.

2. DEFINITION, GENERAL RESULTS, AND FURTHER MOTIVATION

In this section we formally define cursed equilibrium, prove its existence in all finite Bayesian games, and develop some general principles and results. We return at the end of the section to a broader discussion of the motivation for our solution concept, comparison to other approaches to analyzing games, and extensions developed in Appendix A.

Consider a finite Bayesian game, $G = (A_1, \ldots, A_N; T_0, T_1, \ldots, T_N; p; u_1, \ldots, u_N)$, played by players $k \in \{1, \ldots, N\}$. The notation A_k denotes the finite set of Player k's actions, where in a sequential game an action specifies what Player k does at each of her information sets; T_k is the finite set of Player k's "types," each type representing different information that Player k can have. For conceptual and notational ease, we include a set of "nature's types," T_0, and define $T \equiv T_0 \times T_1 \times \cdots \times T_N$ as the set of type profiles and p as the probability distribution over T, which we assume is common to all players. Player k's payoff function $u_k : A \times T \to \mathbb{R}$ depends on all players' actions $A \equiv A_1 \times \cdots \times A_N$ and their types. A (mixed) strategy σ_k for Player k specifies a probability distribution over actions for each type, $\sigma_k : T_k \to \triangle A_k$. Let $\sigma_k(a_k|t_k)$ be the probability that type t_k plays action a_k and let $u \equiv (u_1, \ldots, u_N)$.

The common-prior probability distribution p puts positive weight on each $t_k \in T_k$, and p fully determines the probability distributions $p_k(t_{-k}|t_k)$, Player k's conditional beliefs about the types $T_{-k} \equiv \times_{j \neq k} T_j$ of other players (including nature) given her own type $t_k \in T_k$. Let $A_{-k} \equiv \times_{j \neq 0,k} A_j$ be the set of action profiles for players $j \neq k$ (excluding nature, who takes no action) and let $\sigma_{-k} : T_{-k} \to \bigtimes_{j \neq 0,k} \triangle A_j$ be a strategy of Player k's opponents, where $\sigma_{-k}(a_{-k}|t_{-k})$ is the probability that type $t_{-k} \in T_{-k}$ plays action profile a_{-k} under strategy $\sigma_{-k}(t_{-k})$.

The standard solution concept in such games is Bayesian Nash equilibrium.

DEFINITION 1: A strategy profile σ is a *Bayesian Nash equilibrium* if for each Player k, each type $t_k \in T_k$, and each a_k^* such that $\sigma_k(a_k^*|t_k) > 0$,

$$a_k^* \in \arg\max_{a_k \in A_k} \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \times \sum_{a_{-k} \in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) u_k(a_k, a_{-k}; t_k, t_{-k}).$$

In a Bayesian Nash equilibrium, each player correctly predicts both the probability distribution over the other players' actions and the correlation between the other players' actions and types.

As a first step to defining cursed equilibrium, we define for each type of each player the average strategy of other players, averaged over the other players' types. Formally, for all $t_k \in T_k$, define $\overline{\sigma}_{-k}(\cdot|t_k)$ by $\overline{\sigma}_{-k}(a_{-k}|t_k) \equiv \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \cdot \sigma_{-k}(a_{-k}|t_{-k})$. When Player k is of type t_k, $\overline{\sigma}_{-k}(a_{-k}|t_k)$ is the probability that players $j \neq k$ play action profile a_{-k} when they follow strategy σ_{-k}. A player who (mistakenly) believes that each type profile of the other players plays the same mixed action profile believes that the other players are playing strategies $\overline{\sigma}_{-k}(\cdot|t_k)$ whenever they play $\sigma_{-k}(a_{-k}|t_{-k})$. Note that when players' types are correlated, $\overline{\sigma}_{-k}(a_{-k}|t_k)$ depends on t_k, so different types of Player k have different beliefs about the average action of players $j \neq k$. Let $\overline{\sigma}_{-k}(t_k) : T_{-k} \to \bigtimes_{j \neq 0,k} \triangle A_j$ denote t_k's beliefs about the average strategy of players $j \neq k$, namely $\overline{\sigma}_{-k}(t_k)$ is the strategy players $j \neq k$ would play if each type profile t_{-k} played a_{-k} with probability $\overline{\sigma}_{-k}(a_{-k}|t_k)$.

From this, we define a cursed equilibrium with respect to a parameter $\chi \in [0, 1]$ that measures the degree to which players misperceive the correlation between their opponents' actions and types:

DEFINITION 2: A mixed-strategy profile σ is a *χ-cursed equilibrium* if for each k, $t_k \in T_k$, and each a_k^* such that $\sigma_k(a_k^*|t_k) > 0$,

$$a_k^* \in \arg\max_{a_k \in A_k} \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \times \sum_{a_{-k} \in A_{-k}} \left[\chi \overline{\sigma}_{-k}(a_{-k}|t_k) + (1-\chi)\sigma_{-k}(a_{-k}|t_{-k})\right] \times u_k(a_k, a_{-k}; t_k, t_{-k}).$$

In a χ-cursed equilibrium, each player correctly predicts the probability distribution over her opponents' actions, but she underappreciates the connection between her opponents' equilibrium action profile and their types. Each player plays a best response to beliefs that with probability χ her opponents' actions do not depend on their types, while with probability $1-\chi$ their actions do depend on their types.[5] When $\chi = 0$, χ-cursed equilibrium coincides with Bayesian Nash equilibrium. When $\chi = 1$, each player entirely ignores the correlation between other players' actions and their types. We refer to this extreme case as the *fully cursed equilibrium*, and refer to players in a fully cursed equilibrium as *fully cursed*.

A natural generalization of cursed equilibrium that would help in understanding the theoretical implications of cursed reasoning comes from allowing different players to be cursed to different degrees. While all our theoretical analyses in the remaining sections of the paper assume unitary χ, Appendix A provides definitions along these lines, and much of our empirical analysis below allows for person-specific χ's. While we suspect that many of our basic results for homogenous-cursed equilibrium would extend naturally to heterogenous-cursed equilibrium, the generalization may yield some novel implications.

One important feature of χ-cursed equilibrium—which complicates analysis—is that a player's perception of the strategy played by another player can depend on her own type, and two different players may have different perceptions of the strategy played by a third player. This is impossible in a Bayesian Nash equilibrium, of course, since all types of all players correctly predict the strategies of all types of all other players.[6] When players' types are independent—meaning that for each k, each t_k, t'_k, t_{-k}, $p(t_{-k}|t_k) = p(t_{-k}|t'_k)$—then in any χ-cursed equilibrium each type of Player k as well as any Player j

[5]To see that each player correctly perceives the probability distribution over the other players' actions, note that type t_k of Player k believes that the probability that Players $-k$ play action profile a_{-k} under strategy σ_{-k} is

$$\sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)\big[\chi\overline{\sigma}_{-k}(a_{-k}|t_k) + (1-\chi)\sigma_{-k}(a_{-k}|t_{-k})\big]$$

$$= \chi\overline{\sigma}_{-k}(a_{-k}|t_k) + (1-\chi)\sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)\sigma_{-k}(a_{-k}|t_{-k}) = \overline{\sigma}_{-k}(a_{-k}|t_k).$$

While one could in principle motivate our formal model with a theory that players believe other players play suboptimally given their private information, we do not find that compelling nor is it our motivation. Rather than say that Player A figures out Player B's optimal strategy but believes B does not figure this out, we say that A himself does not properly introspect about how B uses B's private information.

[6]In a Bayesian Nash equilibrium, different players or different types of a given player may have different beliefs about a third player's *actions*, since they may have different beliefs about the likelihood of other players' types. However, by definition, all types of players have common and correct beliefs about others' type-contingent *strategies*. In a cursed equilibrium, different players and types of players may have different beliefs even about these strategies.

share common beliefs about the strategy of any Player $l \neq j, k$. In many of our applications, however, players' types are not independent, so that differences in beliefs prevail in equilibrium.

In many applications, it is both intuitive and convenient to think not in terms of a player's beliefs about others' actions as a function of types, but rather in terms of a player's beliefs about others' types as a function of their actions played. In auctions, for instance, it is often useful to think not in terms of which price each type of bidder bids, but rather which type of bidder bids a given price. Let $\widehat{p}_{t_k}(t_{-k}|a_{-k}, \sigma_{-k})$ be type t_k of Player k's beliefs about the probability of facing type t_{-k} of players $j \neq k$ when they play action profile a_{-k} under strategy σ_{-k}. The following lemma inverts the definition of χ-cursed equilibrium to characterize players' beliefs about other players' types following their actions.

LEMMA 1: *In a χ-cursed equilibrium, for each Player k,*

$$\widehat{p}_{t_k}(t_{-k}|a_{-k}, \sigma_{-k}) = \left((1-\chi)\frac{\sigma_{-k}(a_{-k}|t_{-k})}{\overline{\sigma}_{-k}(a_{-k}|t_k)} + \chi\right)p_k(t_{-k}|t_k). \tag{1}$$

When $\chi = 0$,

$$\widehat{p}_{t_k}(t_{-k}|a_{-k}, \sigma_{-k}) = \frac{\sigma_{-k}(a_{-k}|t_{-k})}{\overline{\sigma}(a_{-k}|t_k)} p_k(t_{-k}|t_k):$$

Player k correctly updates her beliefs about the other players according to Bayes' rule.

When $\chi = 1$,

$$\widehat{p}_{t_k}(t_{-k}|a_{-k}, \sigma_{-k}) = p_k(t_{-k}|t_k):$$

Player k infers nothing about the other players' types from their actions. For intermediate values of $\chi \in (0, 1)$, Player k partially updates to think it more likely that she is facing type t_{-k} when the other players are playing a_{-k}, but she does not fully update.

The following proposition says that χ-cursed equilibria exist in all finite games.

PROPOSITION 1: *If $G = (A, T, p, u)$ is a finite Bayesian game, then for each $\chi \in [0, 1]$, G has a χ-cursed equilibrium.*

The logic behind Proposition 1 is closely related to Lemma 1 and provides a guide for much of our analysis. It is most easily exposited by considering a separating pure-strategy equilibrium, where each type of each player plays a different pure strategy; when t_k observes the action a_{-k} played by types t_{-k}, she believes she is facing t_{-k} with probability $1 - \chi + \chi p_k(t_{-k}|t_k)$ and facing

$t'_{-k} \neq t_{-k}$ with probability $\chi p_k(t'_{-k}|t_k)$. In a cursed equilibrium, Player k plays a best response to these beliefs, which means that she acts as if her payoff from playing action a_k when facing action a_{-k} and type profile t_{-k} is

$$\overline{u}_k^{\chi}(a_k, a_{-k}; t_k, t_{-k}) \equiv (1-\chi)u_k(a_k, a_{-k}; t_k, t_{-k}) + \chi \sum_{\tau_{-k} \in T_{-k}} p_k(\tau_{-k}|t_k) \cdot u_k(a_k, a_{-k}; t_k, \tau_{-k}). \tag{2}$$

This is the χ-weighted average of her actual payoff as a function of actions and types, and her "average" payoff as a function of actions and her own type, averaged over the types of other players. We prove Proposition 1 by noting that since a χ-cursed equilibrium in $G = (A, T, p, u)$ is equivalent to a Bayesian Nash equilibrium in the *χ-virtual game* $\overline{G}^{\chi} \equiv (A, T, p, \overline{u}^{\chi})$, G has a cursed equilibrium whenever $\overline{G}^{\chi}$ has a Bayesian Nash equilibrium. Below we frequently use this reinterpretation and reformulation of cursed equilibrium as the Bayesian Nash equilibrium of $\overline{G}^{\chi}$.

Proposition 1 follows from the fact that whenever G is finite, $\overline{G}^{\chi}$ is finite, and finite games have at least one Bayesian Nash equilibrium. Proposition 1 is of limited general interest, however, since most games we consider have uncountably infinite type and action spaces. Moreover, the existence of a Bayesian Nash equilibrium (which is to say, a $\chi = 0$ cursed equilibrium) is neither necessary nor sufficient for the existence of a χ-cursed equilibrium for each $\chi \in (0, 1]$. However, every game we consider in this paper has an equilibrium for each value of χ, and all existential counterexamples we have devised involve games with discontinuous payoffs or noncompact action spaces. We suspect that cursed equilibria exist in virtually all games of economic interest.

In a cursed equilibrium, a player maximizes her payoffs under the mistaken belief that other players' actions depend less on their types than these actions actually do. We establish in Proposition 2 that if no player can learn anything about her expected payoff from any action profile by learning any other player's type, then the set of cursed equilibria coincides with the set of Bayesian Nash equilibria. To formally state the proposition, we distinguish between the set of Player k's opponents and the set of possible states of the world. Let $T_{-0k} \equiv \times_{i \neq 0,k} T_i$ be the set of possible types of all players $i \neq k$ excluding nature, Player 0. Let $E[U_k(a_k, a_{-k}; t_k, t_{-k})|t_k]$ be Player k's expectation of her payoff when she plays action a_k and the other players play action a_{-k}, conditional on her type t_k; U_k is random because it may depend on t_0 or t_{-0k}. Let $E[U_k(a_k, a_{-k}; t_k, t_{-k})|t_k, t_{-0k}]$ be Player k's expectation of her payoff when she plays action a_k and the other players play action a_{-k}, conditional on her type t_k and the other players' (excluding nature's) type t_{-0k}.

PROPOSITION 2: *If for each Player k, each type $t_k \in T_k$, each type profile $t_{-0k} \in T_{-0k}$, and each action profile $(a_k, a_{-k}) \in A$, $E[U_k(a_k, a_{-k}; t_k, t_{-k})|t_k,$*

$t_{-0k}] = E[U_k(a_k, a_{-k}; t_k, t_{-k})|t_k]$, *then for each* $\chi \in [0, 1]$ *the set of* χ*-cursed equilibria coincides with the set of Bayesian Nash equilibria*.

The condition that $E[U_k(a_k, a_{-k}; t_k, t_{-k})|t_k, t_{-0k}] = E[U_k(a_k, a_{-k}; t_k, t_{-k})|t_k]$ not only requires that no player's payoff be affected by another player's type, but also that no player can learn anything about her expected payoff by learning another player's type. This means essentially that (given a player's type) other players' types are uncorrelated with the state of nature. This distinction is crucial in many of our applications. In a common-values auction, for instance, bidders may not care about other bidders' signals per se, but only about the uncertain value of the object. However, if one bidder learned another bidder's signal, her beliefs about the value of the object and, therefore, about her expected payoff from a profile of bids would change. Hence, Proposition 2 does not apply to common-values auctions, but it does apply to private-values auctions, where each bidder's payoff is a deterministic function of her own type and the profile of bids.

The intuition behind the proposition is that if a player learns nothing about her expected payoff from knowing the other players' types, then misunderstanding the relationship between the other players' types and actions does not matter. If she correctly predicts the probability distribution over the other players' actions, then she will choose the same action irrespective of her theory of which types of the other players play which action.

A final general result is of interest in some applications—and helps give more intuition about the nature of cursed equilibrium. By analogy with pooling equilibria in simple signalling games, say that a strategy profile σ is *pooling* if for each player k there exists some $a_k \in A_k$ such that, for each $t_k \in T_k$, $\sigma(a_k|t_k) = 1$.

PROPOSITION 3: *If a pooling strategy profile* σ *is a* χ*-cursed equilibrium for some* $\chi \in [0, 1]$, *then* σ *is a* χ'*-cursed equilibrium for each* $\chi' \in [0, 1]$.

Proposition 3 implies that every "pooling" Bayesian Nash equilibrium—meaning no player's action depends on her type—is a χ-cursed equilibrium for every value of χ, and any pooling χ-cursed equilibrium is a Bayesian Nash equilibrium. In a pooling equilibrium, players' actions are independent of their types, so ignoring the relationship between others' actions and their information is not a mistake.

Cursed equilibrium combines the assumption that the players have correct beliefs in predicting the *behavior* of other players with the assumption that players have incorrect beliefs in predicting the *strategies* of other players. At first blush, this combination may seem unlikely. Yet it may be justified if each Player A has both a broad sense of Player B's possible types and an empirical sense of Player B's behavior, but does not (fully) think through the logic of the connection between the two. Moreover, without verifying it in a formal model, one learning story seems to justify fully cursed equilibrium. In environments where repeated generations of players (i) have correct a priori beliefs about a game's informational structure, (ii) observe the behavior of a large number of players previously playing the game—but not players' final payoffs—and (iii) are entirely unsophisticated about other players' strategic incentives, then fully cursed equilibrium seems like a natural long-run "equilibrium." Such a dynamic seems plausible, for instance, in some common-value-auction environments. Each new generation of bidders may have access to previous auctions' bids, allowing them to develop a good empirical sense of the distribution of bids without learning either past bidders' information or winning bidders' eventual valuations.[7] Whether one could find a learning story combined with assumptions about a priori partial strategic sophistication that would provide foundations for our exact specification of *partially* cursed equilibrium seems more doubtful.[8]

All said, however, our primary motivation for defining cursed equilibrium is not based on learning or any other foundational justification, but rather on its pragmatic advantages as a powerful empirical tool to parsimoniously explain data in a variety of contexts.

Due to both its different intuitive foundations and its crisply different predictions, cursed equilibrium explains many behavioral departures from Bayesian Nash equilibrium that existing alternatives cannot. The widely used model of quantal-response equilibrium (QRE) developed by McKelvey and Palfrey (1995), for instance, is not well suited to explain the particular patterns of erroneous play we emphasize. On the other hand, it is clear that players make many errors entirely unrelated to cursedness, and we believe that QRE and related approaches are complements rather than substitutes for cursedness. Indeed, like all such analysis, our subsequent empirical analyses of laboratory behavior assumes that there is unexplained random "noise" in the behavior of subjects that our model does not explain. We view this empirical approach as a preliminary step to a more sophisticated QRE-style way to combine unsystematic errors, best responding, and cursedness into a model.

More related is Jehiel's (2005) notion of "analogy-based expectation equilibria," where players bundle other players' decision nodes into "analogy classes";

[7]This interpretation is far more compelling when bidders' types are independent than when correlated, where bidders need to learn the distribution of their opponents' bids conditional on their own type.

[8]Our model may fare no worse than prominent existing models in its lack of learning-theoretic foundations. Game-theoretic research that combines strong but imperfect forms of strategic sophistication with empirical learning is scarce. Indeed, Dekel, Fudenberg, and Levine (2004) have shown with formal models of learning just how difficult it is to justify many seemingly reasonable solution concepts, including the popular approach of Bayesian Nash equilibrium without common priors. The lack of firm learning-theory foundations extends even to Bayesian Nash equilibrium itself. In many environments, correctly learning others' strategies seems far less plausible than learning their behavior.

each player correctly predicts other players' average play within each analogy class but incorrectly assumes that they play the same mixed action profile at each decision node within the class. Analogy-based expectation equilibria nicely captures the idea that players do not distinguish the relationship between others' behavior and information. Jehiel (2005) applies his solution concept to multi-stage games of complete information, where different nodes correspond to different histories of play; his work does not include games of incomplete information. Jehiel and Koessler (2005) have more recently applied this approach to Bayesian games, where an analogy class partitions other players' types. One such partition corresponds to our fully-cursed equilibrium; hence, a special case of their model coincides with a special case of ours.

Another psychologically motivated model of errors recently applied to informational games (and financial markets) is *overconfidence*, which is typically conceptualized as players thinking that they know more than they objectively do or exaggerating their knowledge relative to other players'. While we suspect that overconfidence is psychologically real and economically important, it may be invoked somewhat implausibly in cases where something closer to cursedness drives behavior. The misidentification is natural enough given the previous lack of focus on the type of strategic naivety captured by our model and given the broad-stroke similarity in the two models in many settings of interest. In many games where multiple players have private information, assuming a player overweights his own private information provides a back-door mechanism for capturing his propensity to underweight the informational content in other people's behavior. The two theories' predictions often coincide, for instance, in settings where Bayesian Nash equilibrium predicts "no trade" because participants ought to realize that others' willingness to trade reflects their own private information. Cursedness says that players simply do not appreciate this informational content of others' willingness to trade; overconfidence (depending on how it is modelled) says that people discount others' information because they believe it is weaker than their own. Both can predict overtrading.

Yet the two are very distinct in other settings, even other trade settings. In the simple lemons example in the Introduction, cursedness predicts that a completely uninformed buyer might buy when she should not because she does not realize that the seller only sells low-value products. Her willingness to trade is transparently not due to overconfidence in her valuation of an object, since she does not have any information whose value she can exaggerate. Hence, while there is some scope for interpreting overtrading in settings with two-sided private information as arising from overconfidence, the fact that we observe the same phenomenon in settings with only one-sided private information suggests the interpretation is likely to be wrong.[9]

[9]Even when the mathematical fit is okay, overconfidence has been invoked in some settings where we believe it is simply psychologically implausible. For instance, experimental subjects who

Finally, we believe that cursed equilibrium has some considerable methodological advantages as an approach to modifying Bayesian Nash equilibrium. With some big caveats stressed in the conclusion, it provides an explicit formula for generalizing Bayesian Nash equilibrium that can be applied universally to any existing economic model where available actions are independent of private information. No additional ancillary assumptions are needed and once the parameter χ is chosen, the theory is exactly as parsimonious—and requires exactly the same information—as Bayesian Nash equilibrium.[10] Wherever economists and game theorists currently apply $\chi = 0$ cursed equilibrium, they can also apply $\chi \neq 0$ cursed equilibrium. Although below we estimate different values of χ for different experiments—and sometimes for different players—one eventual test of the model's explanatory power is if a single value of χ (for each person) usefully explains the person's behavior across different games. If it were the case that (say) $\chi = 0.3$ *applied universally* better explains behavior than $\chi = 0$, it would seem to us to be a better model.[11] As it stands now, in fact, *any* value of $\chi > 0$ provides a better fit to the data than does $\chi = 0$ in all the experiments analyzed in this paper with the exception of the jury data reported in Section 5.

The explicit generalization of Bayesian Nash equilibrium should also facilitate careful, normal-scientific estimation and hypothesis testing that is difficult without a precise model. In a few papers, researchers have anticipated our general definition of fully cursed equilibrium in specific contexts and explicitly compared Bayesian Nash equilibrium to fully cursed equilibrium. We are aware of no papers that do what our model enables researchers to do: estimate a degree of cursedness as a parameter in a model that embeds the Bayesian Nash equilibrium as a special case. We believe, for instance, that the lack of a parameterized general model that embeds Bayesian Nash equilibrium has

receive their private signals as draws from an urn are unlikely to think their own signals are more informative than others'. Indeed, insofar as experimental studies where the information structure is not conducive to overconfidence produce the same phenomena as more naturalistic settings, overconfidence seems a less plausible explanation in those naturalistic settings than it at first may appear.

[10]By contrast, Jehiel's (2005) predictions depend crucially on which analogy classes are exogenously specified, as do variants of Jehiel and Koessler's (2005) model that are not equivalent to fully cursed equilibrium. We do not consider this a very strong critique of his model, however. Economists are accustomed to using good psychological and economic intuition to specify informational and payoff assumptions as primitives of games, and we see no reason why they should be unwilling or incapable of similarly developing intuitions for good analogy classes. Nevertheless, because it requires less game-specific judgment than Jehiel's approach, we consider our model more amenable to the simple, universal application of the assumption that players underappreciate the contingent nature of other players' decisions.

[11]The model's explanatory power could also be improved through context- and player-specific parameterization. As we show in some of our data analysis below, for instance, there is empirical evidence for the intuitive notion that with experience players become less cursed. A model that allows for different degrees of cursedness based on experience and other factors could well provide enough additional explanatory power to warrant the decrease in parsimony.

very much hindered empirical investigation of the winner's curse in auctions. At the end of Section 4, we discuss some field studies that clearly hint at the existence of a winner's curse among bidders for oil fields. A common interpretation of these papers is that they demonstrate the nonexistence of the winner's curse in the auctions studied. We believe that this interpretation comes in part from a lack of formal models that allow empirical researchers to directly measure the winner's curse. Differing interpretations could be formally tested with structural models that incorporate cursedness, and would produce point estimates and confidence intervals for χ. More generally, our reading of many empirical papers that study settings with private information is that full empirical testing would yield point estimates for χ that are nontrivially greater than zero.

3. TRADE

In many economic exchanges, one party has private information about the value of the good she might buy or sell that determines the price at which she is willing to trade. This section fleshes out the implications of cursed equilibrium in such settings, with both one-sided and two-sided asymmetric information. We show that trade may occur when Bayesian Nash equilibrium predicts no trade and may not occur when Bayesian Nash equilibrium predicts trade.

We begin by studying one-sided asymmetric information of the sort introduced in Akerlof's (1970) lemons model, which we formalize along the lines Samuelson and Bazerman (1985) formulated in designing an experimental test. A firm offers itself for sale to a raider; the firm knows its book value, but the raider does not. The raider has correct priors that the book value of the firm is uniformly distributed on $[0, 1]$. Whatever its book value, the firm values itself at its book value, while the raider values the firm at $\gamma \geq 1$ times book value. The raider must make the firm an offer, which the firm then accepts or rejects; without loss of generality we take the raider's offer space to be $[0, 1]$. The raider seeks to maximize her expected surplus and the firm accepts any offer above its book value.

Formally, there are two players F (firm) and R (raider), with $T_F \equiv [0, 1]$. The raider, who has no private information, chooses a price $b \in [0, 1]$ at which she offers to buy the firm. The firm chooses a response policy $a : [0, 1] \to \{0, 1\}$, where $a(b) = 1$ means that it accepts the raider's offer of b. The firm's optimal strategy is clear: it sells at price b if and only if (iff) its book value is less than b. Given the uniform distribution of the firm's type, the average value of firms sold at price b is $\frac{b}{2}$, which in turn means the raider's expected surplus from offering b is $b(\gamma\frac{b}{2} - b)$. By familiar "lemons" logic, the lower the bid, the lower the average value of firms that sell. When $\gamma < 2$, the expected net return to the raider will be negative for any positive b, so the unique Bayesian Nash equilibrium outcome is $b = 0$. When $\gamma > 2$, the raider's expected profit is positive whatever her bid and it is maximized at $b = 1$.

What are the χ-cursed equilibria? One is that the firm rejects all bids and the raider offers zero. Yet the best response of some types of firms to a positive offer is to accept. Henceforth we limit our attention to an analog of perfect Bayesian equilibrium, where the firm best responds to some beliefs off the equilibrium path.[12] Consider first the extreme case where $\chi = 1$, so the raider incorrectly thinks that the firm's decision whether to accept the offer does not depend on its book value. Let $\overline{\sigma}_F(a)$ be the average (across types) probability that a firm plays action a. Thus, $\overline{\sigma}_F(1) = \int_0^b 1\,dt + \int_b^1 0\,dt = b$, because firms valued less than b sell while those valued above b do not. In a fully cursed equilibrium, the raider thinks that if she offers b, each firm accepts with probability b. Her perceived payoff from offering b is therefore $b(\frac{\gamma}{2} - b)$, which is maximized by $b = \frac{\gamma}{4}$ for $\gamma \leq 4$ (and at $b = 1$ for $\gamma > 4$). The raider's *true* payoff from bidding $\frac{\gamma}{4}$ is $\frac{\gamma}{4}(\gamma\frac{\gamma}{8} - \frac{\gamma}{4}) = \frac{\gamma^3-2\gamma^2}{32} < 0$ for $\gamma < 2$. Thus the raider suffers a "winner's curse": she does not realize that the firm accepts her offer only when its value is low. The fact that the raider thinks that some firms with values above her bid will sell keeps her from lowering her bid to zero.[13]

For $\gamma \in (2, 4)$, the raider bids too low. Her payoff from bidding $b = \frac{\gamma}{4}$ is $\frac{\gamma^3-2\gamma^2}{32}$, less than the $\frac{\gamma-2}{2}$ from bidding $b = 1$. The predictions of both overbidding when $\gamma < 2$ and underbidding when $\gamma > 2$ come, in fact, for the same reason: a cursed buyer does not fully appreciate the extent to which raising her offer raises the expected value of the goods she buys, and so she pays more attention to how her bid affects her probability of completing a trade than to how it affects the quality of the good she will get.

Now consider $\chi \in (0, 1)$. If the raider offers b, the firm sells iff its valuation is less than b, but in a χ-cursed equilibrium, the raider thinks a firm of type t_F sells with probability

$$(3) \qquad (1-\chi)\sigma_F(1|t_F) + \chi\overline{\sigma}_F(1) = \begin{cases} 1-\chi+\chi b, & \text{for } t_F < b, \\ \chi b, & \text{for } t_F > b. \end{cases}$$

The raider thinks that with probability χ, the firm accepts a bid b with probability b independent of its type, and with probability $1-\chi$, a firm accepts b iff

[12]Formally, σ is a χ perfectly cursed equilibrium of G if it is a perfect Bayesian equilibrium of the χ-virtual game $\overline{G}^{\chi}$. This formulation restricts beliefs off the equilibrium path to be consistent with inference in a cursed equilibrium. A previous version of this paper (Eyster and Rabin (2002)) motivates this refinement at greater length.

[13]Notice that even when $\gamma < 1$, cursed equilibrium involves $b > 0$: even though the raider *knows* that the firm is always worth less to her than to the firm, she still makes a positive offer. Players trade despite common knowledge that there are no gains from trade. We know of no evidence on this prediction, and this degree of error does not seem *entirely* implausible to us, but it does seem somewhat unlikely and indicates a limitation of our approach discussed further in the conclusion.

$t_F < b$. Hence, the raider's perceived expected surplus from bidding b is

$$b(1-\chi+\chi b)\left(\gamma\frac{b}{2}-b\right)+(1-b)\chi b\left(\gamma\frac{b+1}{2}-b\right), \tag{4}$$

which is maximized by $b^* = \frac{\chi\gamma}{4-2\gamma(1-\chi)}$. From this, it can be seen that $\frac{\partial b^*}{\partial\chi} > 0$ if and only if $\gamma < 2$, which means that the buyer overpays when $\gamma < 2$ and underpays when $\gamma > 2$.

Existing experimental evidence on this model shows that subjects do bid positive amounts, contradicting the Bayesian Nash prediction of 0, but in fact they tend to bid in excess of the levels predicted by even the fully cursed equilibrium. When $\gamma = \frac{3}{2}$, the fully cursed equilibrium is $b^* = \frac{3}{8}$. Samuelson and Bazerman (1985) find that the majority of subjects make offers in $(0.5, 0.75)$. Ball, Bazerman, and Carroll (1991) allow subjects to learn by repeating the game 20 times, where subjects learn their payoffs after every round. Such learning does not appreciably affect average bids, which over the course of the trials fall modestly from 0.57 to 0.55.

Holt and Sherman (1994) consider a variant of this model where the raider's priors on the value of the firm are distributed uniformly on $[v_0, v_0 + r]$. Like in Samuelson and Bazerman's model, in Holt and Sherman's model a fully cursed raider can either bid lower than, equal to, or higher than an uncursed raider, depending on the parameter values. For each of the three combinations of γ, v_0, and r that Holt and Sherman tested in laboratory experiments, Table I presents both the χ-cursed equilibrium values of b and the subjects' average bid $\overline{b}$.

Holt and Sherman designed their "no-curse" treatment such that the fully cursed equilibrium coincides with the Bayesian Nash equilibrium; as a result, bids do not depend on χ. In this case, subjects bid quite close to the theoretical prediction. In the winner's-curse treatment, a fully cursed raider would bid 3.56, while an uncursed raider would bid 3. Subjects' average bid was 3.78, slightly above the fully cursed prediction. In the loser's-curse treatment, a fully cursed raider would bid 0.81, while an uncursed raider would bid 1. Subjects'

TABLE I

ADVERSE SELECTION[a]

Curse	r	v_0	γ	$b(\chi)$	$b(\chi=0)$	$b(\chi=1)$	$\overline{b}$
No curse	2	1	1.5	2	2	2	2.03
Winner's	4.5	1.5	1.5	$\frac{45\chi+12}{4+12\chi}$	3	3.56	3.78
Loser's	0.5	0.5	1.5	$\frac{9\chi+4}{4+12\chi}$	1	0.81	0.74

[a] From Holt and Sherman (1994).

average bid was 0.74. Thus, subjects' behavior is much closer to the fully cursed than the Bayesian Nash prediction, although, as in the above treatments, average bids depart too extremely from Bayesian Nash equilibrium to be described adequately by cursed equilibrium. It is clear that fully cursed equilibrium is a much better fit for these lemons data than Bayesian Nash equilibrium, although the overly extreme degree of cursedness that we estimate suggests an aspect of behavior in this setting not captured by our model.

We now turn to two-sided asymmetric information and show that trade can occur in a χ-cursed equilibrium, even when it is common knowledge that the value of the good is identical for the two parties—so that Bayesian Nash equilibrium predicts no trade. While we know of no experimental evidence in such a situation, our prediction of trade matches the common intuition that speculative trade occurs when the no-trade theorems of Milgrom and Stokey (1982) and others predict none. Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ be the set of possible payoff-relevant states of the world, where the two players share the common prior $\mu(\omega_1) = \mu(\omega_2) = \mu(\omega_3) = \frac{1}{3}$. Each trader has private information about the state of the world: Trader A learns when the state is ω_1, but cannot differentiate between states ω_2 and ω_3; Trader B learns when the state is ω_3, but cannot differentiate between states ω_1 and ω_2. The information partitions $\mathcal{P}_A = \{\{\omega_1\}, \{\omega_2, \omega_3\}\}$ and $\mathcal{P}_B = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$ represent Trader A's and B's information, respectively; P_i is an element of Trader i's partition $\mathcal{P}_i$.

A state-contingent trade $t : \Omega \to \mathbb{R}$ specifies a real-valued transfer from Trader A to Trader B in each state of the world. Consider the trade $(1, \varepsilon, -1)$ for $\varepsilon > 0$, meaning that Trader A pays B 1 in ω_1, ε in ω_2, and -1 in ω_3. Trader A would reject this ex ante, as her expected payoff is $-\frac{\varepsilon}{3} < 0$. Now suppose that each trader receives her private information P_i and has the opportunity to reject $(a_i(P_i) = 0)$ or accept the trade $(a_i(P_i) = 1)$, which is implemented if and only if both players consent $(a_A(P_A)a_B(P_B) = 1)$. We claim that

$$a_A(P_A) = \begin{cases} 0, & P_A = \{\omega_1\}, \\ 1, & P_B = \{\omega_2, \omega_3\}, \end{cases} \quad \text{and} \tag{5}$$

$$a_B(P_B) = \begin{cases} 1, & P_B = \{\omega_1, \omega_2\}, \\ 0, & P_B = \{\omega_3\}, \end{cases}$$

constitute a χ-cursed equilibrium for $\varepsilon \le \frac{\chi}{2-\chi}$. First, note that Trader B's strategy is weakly dominant. Clearly Trader A rejects the trade in ω_1. Trader A's expected gain in a χ-cursed equilibrium in states $\{\omega_2, \omega_3\}$ is $\frac{1}{2}[(1-\chi)(-\varepsilon) + (\chi/2)(-\varepsilon + 1)]$, which is positive for $\varepsilon < \frac{\chi}{2-\chi}$. Cursedness leads A to accept the trade in states $\{\omega_2, \omega_3\}$ because she mistakenly believes that B sometimes accepts the trade in state ω_3.

In the example, trade actually occurs with positive probability but not probability 1: it occurs in ω_2 but not in ω_1 or ω_3. This turns out to be a general feature of cursed trade. To see this, consider a finite state space Ω, where two traders share the prior μ over Ω such that for each $\omega \in \Omega$, $\mu(\omega) > 0$. The traders have information partitions $\mathcal{P}_A$ and $\mathcal{P}_B$, and their pure strategies $a_i : \mathcal{P}_A \to \{0, 1\}$ specify whether to accept the trade at every element of their information partitions. A trade $t : \Omega \to \mathbb{R}$ specifies a real-valued transfer from Trader A to Trader B in each state of the world. For simplicity, we impose the tie-breaking assumption that any trader indifferent between accepting and rejecting a trade rejects it.

PROPOSITION 4: *In no χ-cursed equilibrium for any $\chi \in [0, 1]$ is it common knowledge at $\omega \in \Omega$ that $a_A(P_A(\omega))a_B(P_B(\omega)) = 1$.*

Trade cannot be common knowledge in a cursed equilibrium. If it were, then traders would be playing a pooling equilibrium in the element of the meet of $\mathcal{P}_A$ and $\mathcal{P}_B$ that contains ω.[14] If a pooling equilibrium exists on M, then it also exists on the Bayesian game defined by deleting $\Omega \backslash M$ from the original trading game. Proposition 3 would then imply that trade occurs in a Bayesian Nash equilibrium, which would violate Milgrom and Stokey's (1982) no-trade theorem. Thus, while cursed trade may occur with arbitrarily high probability, it can never be commonly known.[15]

4. COMMON-VALUES AUCTIONS

This section illustrates the implications of cursed equilibrium in common-values auctions. We focus on second-price auctions, showing how cursedness leads bidders with certain signals to bid higher than in Bayesian Nash equilibrium and those with other signals to bid lower. We then show that in auctions with a large number of bidders, cursed equilibrium gives rise to the winner's curse: the average winning bid exceeds the average value of the object being auctioned. This implies that in auctions with many bidders, cursedness increases the seller's expected revenue. Yet we provide an example of an auction with only two bidders, where cursedness decreases the seller's expected revenue. Following these theoretical results, we fit cursed equilibrium to a range of laboratory experiments on common-values auctions.

A single object whose value is common to N bidders is to be sold by auction. Let $x = (x_1, \dots, x_N) \in [\underline{x}, \overline{x}]^N \subset \mathbb{R}^N$ be a profile of the bidders' real-valued

[14] The element of the meet of $\mathcal{P}_A$ and $\mathcal{P}_B$ that contains ω is the smallest set $M \ni \omega$ such that for each $i \in \{A, B\}$, $P_i \in \mathcal{P}_i$, $M \cap P_i = \emptyset$, or $P_i \subset M$. It contains all states that A thinks possible, that A thinks B thinks possible, that A thinks B thinks A thinks possible, etc., and vice versa.

[15] Trade can, however, be known, as the following example illustrates. Using the same information partitions and priors as in the example above, consider the trade that calls for B to pay A 1 in state ω_1, -2 in ω_2, and 3 in ω_3. It is a fully cursed equilibrium for Trader A to always accept, and B to reject in ω_3 and accept otherwise. In ω_1, B knows he will accept, and since A always accepts, B knows there is trade. Trader A knows the state is ω_1 and, therefore, that $P_B = \{\omega_1, \omega_2\}$ and $a_B(P_B) = 1$, so she too knows that there is trade. Trade is not commonly known at ω_1 as B does not know that A knows they trade.

private signals about the value of the object being auctioned and let $s \in \mathbb{R}$ be a characteristic of the object.[16] The value of the object, $u(x, s) \geq 0$, is a continuous and increasing function of bidders' signals, x, and of the characteristic, s; we assume u to be symmetric in the x_i's. The experimental literature focuses on two cases: $u(x, s) = \sum_{i=1}^{N} x_i$ and $u(x, s) = s$. We assume that the $N + 1$ random variables $S, X_1, \ldots, X_N$ have density $f(x, s) = \prod_{i=1}^{N} g(x_i|s)h(s)$, where $g(x_i|s)$ satisfies the monotone-likelihood ratio property: for $x_i' \geq x, s' \geq s$,

$$\frac{g(x_i'|s')}{g(x_i|s')} \geq \frac{g(x_i'|s)}{g(x_i|s)}.$$

Each bidder is risk neutral and her utility from winning the auction is the common value of the object, $u(x, s)$ less the price she pays, p; her utility from losing the auction is zero. To analyze cursed equilibrium in common-values auctions, we use the χ-virtual game introduced in Section 2, where Bidder i's utility from winning the auction at price p when the value of the object is $u(x, s)$ is $(1 - \chi)u(x, s) + \chi E[u(X, S)|X_i = x_i] - p$, where x_i is Bidder i's signal about the value of the object. Bidder i's valuation of the object is the χ-weighted average of the object's actual value and her expectation of its value conditional on her signal. A Bayesian Nash equilibrium in the auction where bidders' have these utilities is a cursed equilibrium of the original auction.

Let $Y_{-i}(1) \equiv \max_{j \neq i} X_j$ be the highest signal received by a bidder other than Bidder i. Two functions that play important roles in our analysis merit definition here: $r(x_i) \equiv E[u(X, S)|X_i = x_i]$ is the expected value of the object conditional on Bidder i's receiving the signal x_i, and $v(x_i, y) \equiv E[u(X, S)|X_i = x_i, Y_{-i}(1) = y]$ is the expected value of the object conditional on Bidder i's receiving the signal x_i and the highest of the other bidders' signals being y.

In a second-price auction, the highest bidder wins the auction and pays the second-highest bid. Milgrom and Weber (1982) show that a Bayesian Nash equilibrium of the second-price auction in this setting is $b_i(x_i) = v(x_i, x_i)$—Bidder i bids her expectation of the value of the object conditional on both her signal and the highest of the other bidders' signals being x_i.[17] Intuitively, if Bidders $-i$ follow their equilibrium strategies, the only change in payoff to

[16]Throughout this section, we use uppercase letters to denote random variables and lowercase letters to denote values that these random variables take on.

[17]To see that this is an equilibrium, suppose that bidders $j \neq i$ follow their proposed equilibrium strategies, in which case a Bidder i with signal x_i who bids b_i receives an expected payoff of

$$\int_{\underline{x}}^{b_j^{-1}(b_i)} [v(x_i, y) - v(y, y)] f_{Y_{-i}(1)}(y|X_i = x_i)\, dy,$$

where $f_{Y_{-i}(1)}(\cdot|X_i = x_i)$ is the density of $Y_{-i}(1)$ conditional on $X_i = x_i$. Bidder i's bid b_i only affects the limits of integration, and since $v(x_i, y) \geq v(y, y)$ if and only if $y \leq x_i$ by affiliation, Bidder i's expected utility is maximized when $b_j^{-1}(b_i) = x_i$ or $b_i = b_j(x_i)$.

Bidder i from raising her bid above $v(x_i, x_i)$ comes from winning some auctions where $Y_{-i}(1) > x_i$ and so $v(x_i, y_{-i}(1)) < v(y_{-i}(1), y_{-i}(1))$, namely auctions where the price exceeds the expected value of the object.

In the χ-virtual game corresponding to the second-price auction, Bidder i's expectation of the value of the object conditional on her signal being x_i and the highest of the other bidders' signals being y is

$$(6) \qquad E\big[(1-\chi)u(X,S) + \chi E[u(X,S)|X_i = x_i]|X_i = x_i, Y_{-i}(1) = y\big] \\ = (1-\chi)v(x_i, y) + \chi r(x_i).$$

When $\chi = 1$, bidders act as if the auction were one of private (but uncertain and correlated) values. For $\chi \in (0, 1)$, the χ-virtual game describes an auction where bidders' valuations have both a private and a common component.

PROPOSITION 5: *For each $\chi \in [0, 1]$, it is a χ-cursed equilibrium of the second-price auction for each Bidder i with signal x_i to bid $b_i = (1-\chi)v(x_i, x_i) + \chi r(x_i)$.*

In a χ-cursed equilibrium, rather than bid her expectation of the value of the object conditional on her signal being both the highest and second highest, Bidder i bids the χ-weighted average of that and her expectation of the value of the object conditional on her signal alone. When $\chi = 1$, bidders act as if winning the auction conveys no information about the value of the object; they bid as if in an auction of private (albeit uncertain) values. In general, the second part of Bidder i's bidding function reflects the fact that she underappreciates the information content in winning. Because $r(x_i)$ and $v(x_i, y)$ both increase in x_i, Milgrom and Weber's (1982) argument above establishes that $b_i(x_i) = (1-\chi)v(x_i, x_i) + \chi r(x_i)$ is a Bayesian Nash equilibrium of the χ-virtual auction and, therefore, a χ-cursed equilibrium of the second-price auction.

In the symmetric χ-cursed equilibrium, bidders with high signals overbid relative to Bayesian Nash equilibrium, while those with low signals underbid. This follows from the fact that $\frac{\partial}{\partial \chi} b_i = r(x_i) - v(x_i, x_i)$, which is positive for low x_i—the expected value of the object conditional on one low signal exceeds that conditional on N low signals—and negative for high x_i—the expected value of the object conditional on one high signal is lower than that conditional on two high signals. Consequently, cursedness produces an ambiguous effect on the seller's expected revenue.

We say that Bidder i suffers the winner's curse in equilibrium of a given auction if her expected surplus from entering the auction is negative; that is, if the expectation of the value of the object less the price, both conditional on the event that she wins, is negative. Let $p_i(b)$ denote the price that Bidder i pays when she wins the auction and bids are b; for example, in a first-price auction, $p_i(b) = b_i$.

DEFINITION 3: Bidder i suffers the *winner's curse* in equilibrium (b_i, b_{-i}) if $E[u(X,S) - p_i(b(X))|b_i(X_i) > \max_{j\neq i} b_j(X_j)] < 0$.

In a symmetric equilibrium of a symmetric model where the object always sells, Bidder i suffers the winner's curse if $E[u(X,S)] < E[p(b(X))]$, namely the expected price exceeds the expected value of the object.[18]

Let $(A^N)_{N=2}^{\infty} = (X^N, S, u^N)_{N=2}^{\infty}$ be a sequence of N-bidder, second-price, common-values auctions satisfying all the assumptions in this section and the additional assumption that the densities $g(x_i|s)$ and $h(s)$ do not depend on N. The mapping from bidders' signals and the object's characteristic to its common value, u^N, may depend on N.

DEFINITION 4: A sequence of auctions (A^N) is *competitive* if

$$\lim_{N\to\infty}\big(E[u^N(X^N,S)] - E\big[v^N\big(Y^N(2), Y^N(2)\big)\big]\big) = 0. \tag{7}$$

A sequence of second-price auctions is competitive if the expected price in the Bayesian Nash equilibrium converges to the value of the object. Any sequence comprised of either of the two forms of auction prevalent in the experimental literature (and, indeed, many others) is competitive. Let $\overline{x}(s) \equiv \sup\{x : g(x|s) > 0\}$ be the highest possible signal given the object's characteristic, s.

PROPOSITION 6: *Suppose that the sequence of auctions (A^N) is competitive and that there exists some $\varepsilon > 0$ such that for each N, $E[r^N(\bar{x}(S))] > E[u^N(X,S)] + \varepsilon$. Then for each $\chi > 0$, every bidder in the symmetric χ-cursed equilibrium of the N-bidder, second-price auction A^N suffers the winner's curse when N is sufficiently large.*

The assumption in the proposition rules out cases where bidders' signals provide no meaningful information about the object's value or become sufficiently uninformative as the number of bidders grows. Both forms of auction studied in the experimental literature satisfy this assumption.

Proposition 6 establishes that whenever cursed bidders are uncertain about the value of the object, then as long as they are cursed—no matter how slightly—they suffer the winner's curse in auctions with sufficiently many bidders. The intuition is quite simple: in competitive auctions, the price converges to the average value of the object in Bayesian Nash equilibrium; with enough

[18]A more liberal definition of the winner's curse would include situations in which a bidder's expected surplus from entering the auction is less than Nash-equilibrium analysis suggests. We choose our definition to emphasize the severity of overbidding and correspond to the folk wisdom that winning bids in common-values auctions tend to exceed the values of the objects being auctioned.

bidders, the pivotal bidder in a cursed equilibrium overbids relative to Bayesian Nash equilibrium and, hence, the price exceeds the average value of the object. Consequently, bidders suffer the winner's curse. Whenever bidders suffer the winner's curse, Milgrom and Weber's (1982) famous linkage principal fails: a seller who could commit to a policy of truthfully revealing (x, s) would prefer not to do so because it would stop cursed bidders from suffering the winner's curse, decreasing expected revenue.[19]

In auctions with many bidders, bidders suffer the winner's curse and, hence, cursedness increases the seller's expected revenue. In auctions with only a few bidders, however, cursedness turns out not always to increase the seller's expected revenue.

EXAMPLE: Let $N = 2$, $X_i \sim U[0, 1]$, and $u(X) = \min\{X_1, X_2\}$. In this auction, two bidders have independent and identically distributed signals, and the value of the object is the lower of their signals. Observe that $v(x_i, x_i) = E[\min\{X_1, X_2\}|X_1 = X_2 = x_i] = x_i$ and $r(x_i) = E[\min\{X_1, X_2\}|X_i = x_i] < x_i$ for all $x_i > 0$. Thus, bidders bid strictly lower for all signal values (other than 0) in a cursed equilibrium than in a Bayesian Nash equilibrium. Cursedness decreases the seller's revenue.

The next proposition stresses the importance of having only two bidders for cursedness to lower the seller's expected revenue.

PROPOSITION 7: *Suppose that $r^N(X)$ is a symmetric random variable and $N \geq 3$. Then the seller's expected revenue in the symmetric χ-cursed equilibrium of the second-price auction increases in χ.*

We do not believe Proposition 7's assumption of symmetry in the signal structure to be an economically meaningful concept. We feel, however, that the result is of interest for two reasons. First, it is a condition satisfied in most of the experimental literature. For example, when $u^N = \sum_{i=1}^N x_i$ and signals are independent, $r^N(x_i) = x_i + (N-1)E[X_i]$; since the random variable $r^N(X_i)$ just rescales X_i, it is symmetric whenever X_i is symmetric. Second, while we believe that for certain asymmetric signal structures the seller's expected revenue may not increase in χ for three bidders, symmetry provides a rough guide as to which auctions can give rise to the winner's curse and which cannot.

PROPOSITION 8: *Suppose that $r(X)$ is a symmetric random variable and that bidders in the symmetric χ-cursed equilibrium of the N-bidder, second-price auction suffer the winner's curse. Then $N \geq 4$.*

[19]The working-paper version of our paper (Eyster and Rabin (2002)) provides an example where a seller who can commit to reveal information about the object's characteristic s (but not about bidders' signals) prefers not to reveal that information to cursed bidders.

A necessary condition for the winner's curse to emerge in a symmetric auction model is that there be at least four bidders. In their overview of the experimental literature on common-values auctions, Kagel and Levin (2002) find that the winner's curse appears only in settings with five or more bidders. While some researchers discussing this pattern have suggested that the psychology of the winner's curse may be qualitatively different in larger auctions than in smaller ones, our model predicts precisely this pattern. For any χ, the winner's curse emerges in a large enough auction, which, in typical lab experiments, requires the presence of no fewer than four bidders.

While we have focused on second-price auctions, it is straightforward to analyze cursed equilibria in first-price and English auctions using bidders' χ-virtual valuations. Since these χ-virtual valuations fit into Milgrom and Weber's (1982) general symmetric model (where a bidder's valuation may depend upon her own signal in a different way than it depends upon other bidders' signals, as long all bidders do this in the same way), their revenue ranking carries through to cursed equilibrium: the auctioneer's expected revenue in an English auction exceeds that in a second-price auction, which in turn exceeds that in a first-price auction.[20]

Rather than explore more general implications of cursed equilibrium in auctions, we conclude this section by using our model to address some of the large body of experimental evidence on common-values auctions. In an early experiment, Bazerman and Samuelson (1983) auctioned off translucent jars of coins to student subjects. Subjects could see the jars but did not know how many coins they contained. The highest bidder paid her bid and received the paper-dollar equivalent of the coins in the jar. In addition, subjects guessed the value of the jars and the subject who guessed closest to the true value won a cash prize. Whereas all jars actually contained \$8.00, the average winning bid was \$10.01. Subjects estimated the jars to be worth only \$5.13 on average. Despite being on average too pessimistic about the value of the jars, subjects suffered the winner's curse, presumably because those with high bids bid close to their estimates, rather than tempering their bids to incorporate the information content in winning.

Avery and Kagel (1997) report experimental evidence on a simple second-price, common-values auction with two bidders. Each bidder receives a signal drawn from a uniform distribution on $[1, 4]$ and the value of the object being auctioned is the sum of the two bidders' signals, $u(x_1, x_2) = x_1 + x_2$. Proposition 5 establishes that the symmetric χ-cursed equilibrium in this auction has

[20] Levin, Kagel, and Richard (1996) provide experimental evidence that English auctions raise less revenue than first-price auctions. While their result is not captured by our formal solution concept, we judge it consonant with our underlying motivation: bidders may better understand the connection between other bidders' bids and private information when they see those bids explicitly, as they do in the English auction, than when they condition on them ex ante. We return to this issue in the conclusion.

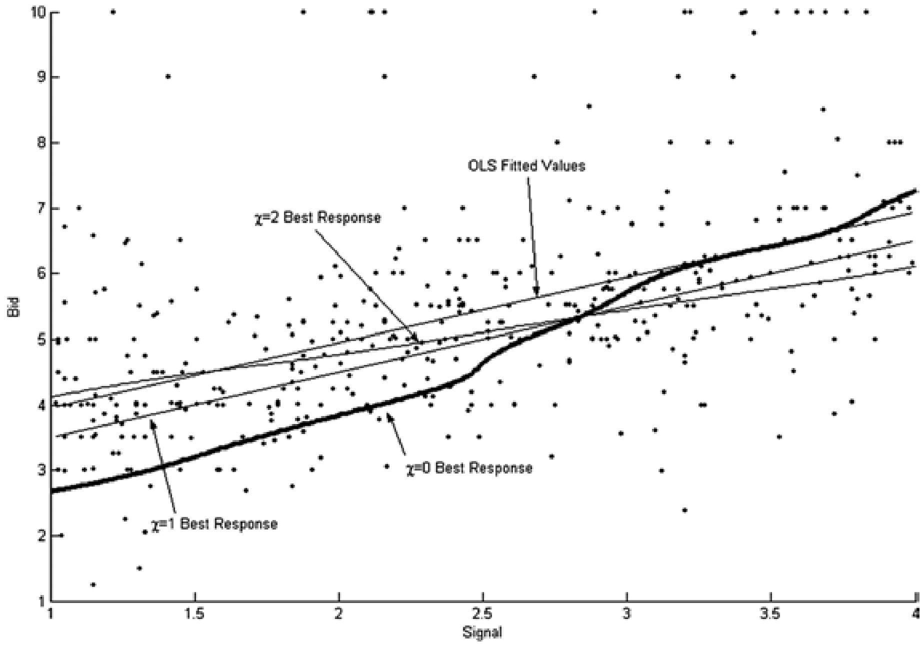

FIGURE 1.—Inexperienced bidders.

Bidder i bid

$$b_i(x_i) = (1-\chi)v(x_i, x_i) + \chi r(x_i) = (2-\chi)x_i + \frac{5\chi}{2}. \tag{8}$$

When $\chi = 0$, bidders bid twice their signals, while when $\chi = 1$, they bid once their signals plus the expected value of the other bidder's signal: cursedness raises the intercept and lowers the slope of a linear bidding function.

Avery and Kagel divide their subjects—mostly undergraduate economics students—into two groups. Inexperienced subjects have already participated in seven (unreported) practice auctions; the reported data cover their next 18 auctions. Experienced subjects are formerly inexperienced subjects with 25 auctions' experience; the reported data cover their next 24 auctions. Figure 1 shows Avery and Kagel's data for inexperienced subjects. The overall pattern of bidding appears consistent with fully cursed equilibrium—depicted by the dotted line—which predicts $b_i = 2.5 + x_i$: subjects with signals below 2.5 bid more than their Nash equilibrium bids of twice their signals, while those with signals above 2.5 bid less than their Nash equilibrium bids. Regressing bids on signals and a constant term gives

$$b_i = \underset{(0.20)}{2.95} + \underset{(0.079)}{1.00} \cdot x_i + \varepsilon_i, \tag{9}$$

where standard errors are given in parentheses. Note, however, that the estimated intercept term is too high to be explained by cursedness alone, as some subjects bid more than the maximum possible value of the object conditional on their signals ($b_i > x_i + 4$), a weakly dominated strategy.

Subjects' bids show substantial variation around the fully cursed equilibrium. Bidders appear not only to be cursed, but also to make errors unexplained by cursed equilibrium (e.g., bidding weakly dominated bids). The broad curve in Figure 1 depicts the best response to subjects' actual strategies. To compute this, we first estimated nonparametrically the expected signal of a bidder bidding b for every possible b. Second, we used this estimate to compute the bid that maximizes a subject's expected payoff as a function of her signal. Figure 1 shows that taking into account the fact that the other bidder does not bid according to Bayesian Nash equilibrium, bidders with low signals bid higher than optimal, while those with high signals bid lower than optimal.

To explore individual-level bids, we enrich our basic model in two ways. First, we allow for the possibility that bidders are heterogeneously cursed, a formal model of which can be found in Appendix A. Second, we allow for the possibility that bidders make other errors uncorrelated with cursedness. Hence, each bidder bids a noisy and cursed best response to the empirical distribution of other bidders' bids and signals. Given the empirical distribution of other bidders' bids and signals, (b_{-i}, x_{-i}), Bidder i with signal x_i and cursedness parameter χ_i maximizes her perceived expected surplus by bidding

$$\begin{aligned}(10)\qquad \widehat{b_i} &= x_i + (1-\chi_i)E[X_j|b_j(X_j) = \widehat{b_i}] + \chi_i E[X_j] \\ &= x_i + (1-\chi_i)E[X_j|b_j(X_j) = \widehat{b_i}] + \frac{5\chi_i}{2}.\end{aligned}$$

(Note that if Bidder j were to bid in the same way, then $E[X_j|b_j(X_j) = \widehat{b_i}] = x_i$ and so $\widehat{b_i} = (2-\chi_i)x_i + \frac{5\chi_i}{2}$, the χ_i-cursed equilibrium of this auction.) We assume that Bidder i actually bids

$$(11)\qquad b_i = \widehat{b_i} + \varepsilon_i = x_i + (1-\chi_i)E[X_j|b_j = \widehat{b_i}] + \chi_i E[X_j] + \varepsilon_i,$$

her intended bid plus a normally distributed error term ε_i that has zero mean and finite variance. To estimate this probability model, we first estimate $E[X_j|b_j = \widehat{b_i}]$ from the data nonparameterically.[21] For each value of χ_i, each signal x_i leads to a fixed point $\widehat{b_i}$; the actual bid, b_i, less $\widehat{b_i}$ yields the error

[21]The auctions for inexperienced bidders were run in two groups, one with 11 bidders and the other with 12 bidders. We have merged these two to improve the quality of our estimate of $E[X_j|b_j = \widehat{b_i}]$. Because several bidders bid high for low signals, our estimates $E[X_j|b_j = \widehat{b_i}]$ decrease for high bids, in which case the fixed point $\widehat{b_i}$ is not unique. Excluding that portion of our data would not substatially change any of the results.

TABLE II

BIDDING IN AVERY AND KAGEL (1997)

Estimated χ's with Standard Errors (SE)				
	Inexperienced Subjects		Experienced Subjects	
Subject	χ	SE	χ	SE
1	2.03	0.22	0.24	0.09
2	0.01	0.35	0.61	0.04
3	3.78	3.87	1.19	0.04
4	1.16	0.17	1.05	0.01
5	0.96	0.02	0.78	0.02
6	0.50	0.17	0.65	0.02
7	2.13	0.89	0.48	0.11
8	1.28	0.07	1.69	0.24
9	1.10	0.02	1.00	0.00
10	5.68	7.86	0.49	0.07
11	6.89	16.00	0.46	0.09
12	3.79	3.46	−0.05	0.20
13	1.85	0.65	−0.35	0.22
14	666.08	343,488.42	15.68	52.45
15	1.51	0.18		
16	1.07	0.02		
17	0.83	0.13		
18	1.43	0.25		
19	819.68	316,873.69		
20	−0.47	0.54		
21	−0.44	0.60		
22	2.30	1.52		
23	3.81	4.54		
All	1.48	0.08	0.63	0.02

term, ε_i. We choose the value of χ_i that minimizes the sum of squared error terms, the nonlinear least squares estimate of χ_i. The second and third columns of Table II provide point estimates of χ_i and their standard errors.

Inexperienced subjects clearly vary in the degree of their cursedness. Most have point estimates near unity or above, and only a small minority have confidence intervals that include zero (best responding to correct beliefs), but the 95% confidence interval of $(1.32, 1.64)$ is too high to be well described by cursed equilibrium, where χ lies between 0 and 1. Indeed, many bids are weakly dominated ($b_i > x_i + 4$), an anomaly not captured by cursedness. Figure 1, by graphing the $\chi = 2$ best response to actual bids, shows how a $\chi > 1$ better fits these weakly dominated bids for low signal holders than does $\chi = 1$.

For a cleaner test, we turn to experienced subjects, who bid weakly dominated bids less frequently. Figure 2 shows the distribution of their bids, which overall exhibits less heterogeneity than that of inexperienced bidders. The re-

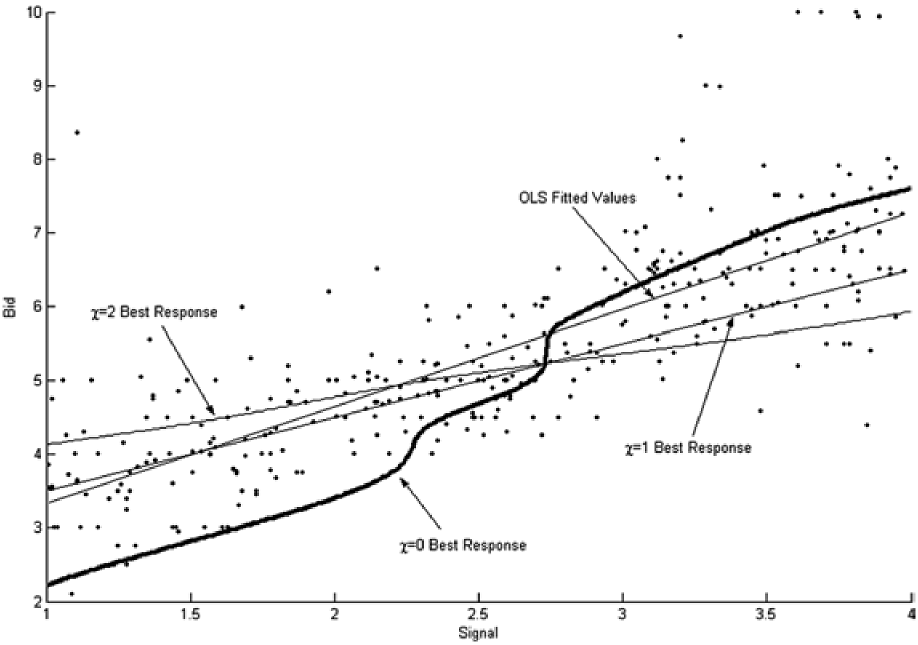

FIGURE 2.—Experienced bidders.

gression line of

$$b_i = \underset{(0.15)}{2.02} + \underset{(0.053)}{1.31} \cdot x_i + \varepsilon_i \qquad (12)$$

is depicted by a dashed line. Its lower intercept and higher slope than that for inexperienced subjects shows that experienced subjects bid closer to Bayesian Nash equilibrium (with an intercept of 0 and slope of 2). Again, the broad curve indicates the best response to the subjects' empirical strategies. Like inexperienced bidders, experienced bidders with low signals bid above their best response, while those with high signals bid below, just the pattern of bidding predicted by cursedness.

Columns 4 and 5 of Table II present estimates and confidence intervals for subject-specific χ's. A $\chi = 0.63$ fits the data best, with a 95% confidence interval of $(0.59, 0.67)$. Only three subjects have estimated confidence intervals that contain best responses to other subjects' play ($\chi = 0$). With the exception of subject 14 (one of those three), who consistently bid above the maximum possible value of the object, experienced subjects' behavior appears to be well described by cursed equilibrium with a value of χ closer to 1 than 0.

To compare how well cursed equilibrium fits the data relative to Bayesian Nash equilibrium, we check the latter's fit against χ-cursed equilibrium for

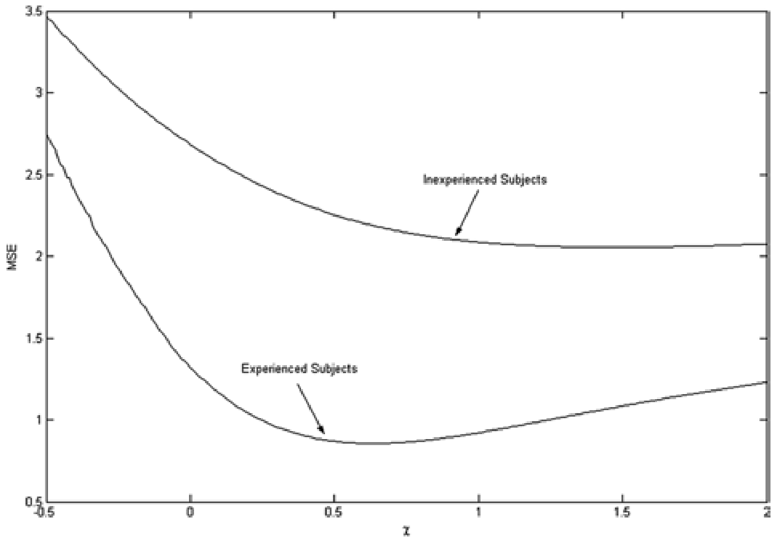

FIGURE 3.—Mean squared errors.

every value of $\chi \in [-0.5, 2]$.[22] To do this, we assign every subject the same value of χ and then use their actual bids to compute the residual sum of squares. Figure 3 shows the residual sum of squares for each value of χ, for inexperienced and experienced subjects, respectively. Increasing χ from -0.5 improves fit for inexperienced subjects until 1.48, our estimate of χ for inexperienced subjects. While cursedness cannot adequately fit the data for the reasons described above, it does fit better than Bayesian Nash equilibrium: any $\chi \in (0, 1]$ fits the data better than $\chi = 0$. Furthermore, the improvement in fit moving from $\chi = 1$ to $\chi = 1.48$ appears quite small relative to that moving from $\chi = 0$ to $\chi = 1$. For experienced bidders, increasing χ from -0.5 improves fit up until $\chi = 0.63$. Here too any $\chi \in (0, 1]$ fits the data better than $\chi = 0$. In sum, not only does the best-fitting χ-cursed equilibrium fit the data better than Bayesian Nash equilibrium, but any χ-cursed equilibrium fits better.

Avery and Kagel's (1997) experimental design allows us to dismiss several potential alternative hypotheses that explain the winner's curse in common-values auctions. For example, bidders who are overconfident about the quality of their private information might bid too close to their expectation of the value of the object conditional on their private information alone, leading to a winner's curse. In Avery and Kagel's (1997) auctions, however, bidders cannot

[22]We chose the interval $[-0.5, 2]$ before conducting the analysis to check whether values of χ outside the allowed parameter space fit better than those inside. Bayesian Nash equilibria exist in the χ-virtual game for any $\chi \in \mathbb{R}$, even though for $\chi \notin [0, 1]$ they are not interpretable as cursed equilibria.

be overconfident in their private information: the object's value is the sum of two halves, where each bidder knows precisely the value of her half. Another explanation for overbidding is that bidders do not share common priors, but it is straightforward to show that Bayesian Nash equilibrium in this auction does not depend on the distribution of bidders' signals and, therefore, noncommon priors about signals would not affect the Bayesian Nash equilibrium.[23] In addition, mistaken beliefs about the distribution of the other bidder's bids cannot explain the pattern of bidding—as long as each bidder understands the other bidder's mapping from signals to bids—by exactly the same argument; neither can risk aversion, which does not affect Bayesian Nash equilibrium in this auction; nor can McKelvey and Palfrey's (1995) quantal-response equilibrium, where players best respond with error to other players' behavior, and the probability that a bidder makes any given error decreases in its equilibrium cost. In this model, bidders with any given signal would bid both above and below their best response to other bidders' bidding rules, in contrast to the pattern observed here.

One alternative hypothesis for Avery and Kagel's results, however, seems plausible. As in many experiments, subjects may have a tendency to choose actions toward the middle of their action spaces, in this case toward the middle of [2, 8], the range of possible values of the object and undominated bids. Such a tendency would lead subjects with signals less than 2.5 to bid higher than in Bayesian Nash equilibrium and those with signals above 2.5 to bid lower. While this hypothesis fits this experiment, it cannot account for anomalous bidding in most of the experimental common-values auction literature, where Nash equilibrium requires that bidders shade their bids below their expectations of the value of the object conditional on their signals by a constant amount. In Eyster and Rabin (2002), we show how cursed equilibrium captures the empirical tendency in these auctions for bidders of all signal levels to bid higher than Bayesian Nash equilibrium predicts.

Kagel and Levin (1986) test a model where the common value of the object, s, is distributed uniformly over $[\underline{s}, \overline{s}]$ and each bidder i receives a signal $X_i \sim U[s - \frac{a}{2}, s + \frac{a}{2}]$. The χ-cursed equilibrium of the first-price auction is

$$b(x_i) = x_i - \frac{a}{2} + \chi a \frac{n-2}{2n} + \frac{a\left(1 - \frac{n-1}{n}\chi\right)}{n+1} z_i, \tag{13}$$

where $x_i \in [\underline{s} + \frac{a}{2}, \overline{s} - \frac{a}{2}]$ and $z_i = \exp(-\frac{n(x_i - (\underline{s} + \frac{a}{2}))}{a})$.[24] For the range of parame-

[23] In any such auction where Bidder 1 has signal X_1 with distribution F_1, Bidder 2 has signal X_2 with distribution F_2, and the value of the object is $X_1 + X_2$, it is a symmetric Bayesian Nash equilibrium for each bidder to bid twice her signal.

[24] Note that the χ-cursed equilibrium in this auction is equivalent to a Bayesian Nash equilibrium of the first-price auction where bidders' have their χ-virtual payoffs. For an explicit derivation, see Eyster and Rabin (2002).

TABLE III

COMMON-VALUES AUCTIONS[a]

n	$a/2$	Observations	χ	Standard Error	Average Winner's Profit
3	18	27	1.59	0.26	4.62
3	24	15	2.03	0.56	5.36
3	30	18	1.35	0.21	5.07
4	12	132	0.42	0.07	2.75
4	18	120	0.41	0.16	4.15
4	30	64	0.89	0.15	8.75
5	12	100	0.38	0.06	2.39
5	18	15	0.60	0.81	−0.31
6	12	12	0.68	0.38	−3.46
6	18	60	0.51	0.13	1.41
6	24	30	0.88	0.14	4.14
6	30	96	0.64	0.13	0.71
7	12	131	0.43	0.09	−1.41
7	18	77	0.70	0.09	−1.37
7	30	56	0.56	0.14	2.06

[a]From Kagel and Levin (1986).

ter values used in the experiments, z_i is negligible and we ignore it henceforth. Table III presents estimates of χ for several of Kagel and Levin's experiments that differ in the number of bidders as well as the noisiness of their signals. These come from simply regressing bids on signals. Hence, we are estimating which χ-cursed equilibrium fits observed bids closest, and not using our cursed-expectations concept, to make the point that even without accounting for other errors that bidders might make, they appear to be significantly cursed. In all treatments, $\frac{db(x_i)}{dx_i}$ is essentially 1, as predicted by χ-cursed equilibrium for any χ.

Cursed equilibrium does not capture three-bidder auctions well, where bidders systematically bid above even fully-cursed-equilibrium bids: all estimates of χ are well above unity. Moreover, bidders in first-price auctions overbid even with private values (see, e.g., Kagel (1995)) and our estimates of χ do not take into account uncorrelated errors that bidders might make. For all other cases, however, χ ranges roughly in $(0.4, 0.9)$. Positive values of χ fit the data much better than does Bayesian Nash equilibrium. Our model of cursedness also help to clarify some interpretations in the auction literature about how the number of bidders affects bidding. Our conceptualization of the source of the winner's curse predicts that profits ought to drop as the number of bidders increases, as empirically observed, without any switch in underlying psychology. Table III in fact shows that, while average profits drop as the number of bidders increases, there is no discernible pattern of χ increasing. For instance, by the crude measure of taking the average χ among auctions with the same number of players weighted by number of players, the average χ

is 1.63, 0.51, 0.41, 0.64, and 0.54 for 3, 4, 5, 6, and 7 bidders, respectively. Tracing out the estimated χ as a function of number of players for each value of $\frac{a}{2}$ likewise does not suggest an increase in χ.

Finally, an empirical literature examines the winner's curse in field data. In a series of papers, Kenneth Hendricks, Robert Porter, and others have examined bidding in auctions for oil and gas leases. Bidders' valuations in these auctions include a large common component about which they have private information. Hendricks, Porter, and Boudreau (1987) calculate ex post profits by combining actual production and oil prices with the winning bids and estimates of the costs of production. While firms make positive profits on average, Hendricks, Porter, and Boudreau (1987) test whether they bid optimally by checking whether scaling of a firm's bids by a constant factor would increase profits, a test substantially weaker than optimality. Thirteen of their eighteen firms (submitting 77% of total bids) would have increased profits by reducing their bids by a constant factor. In auctions where bidders have roughly symmetric information, while net profits are positive overall, they are negative in auctions with seven or more bidders. In a high-stakes environment with big firms that win many auctions, winning bids appear to be too high.

Hendricks and Porter (1988) examine oil auctions where some bidders own adjacent tracts (neighbors) that presumably make them better informed than their competitors without neighboring tracts (nonneighbors). If nonneighbors have only noisy versions of neighbors' information and are uncertain who bids, then they should lose money on auctions where no neighbors bid and win money on auctions where neighbors bid. In fact, nonneighbors lost an average of \$2.69 million (with a standard deviation of \$0.86 million) in auctions where no neighbors bid and made an average of \$0.78 million (with a standard deviation of \$2.64 million) in auctions where at least one neighbor bid. Overall, nonneighbors lost an average of \$0.42 million (with a standard deviation of \$1.76 million). While the negative and insignificant point estimate on overall profits cannot conclusively pinpoint the winner's curse, it does suggest that winning bids may have been too high. To the extent that nonneighbor bidders had any private information about the value of the oil, paid any costs to enter the auction, etc., they should have earned strictly positive expected profits. The fact that they did not suggests that a structural model designed to empirically estimate χ in this setting would yield a $\chi > 0$.

5. INFORMATION REVELATION

In this section, we consider a range of further applications of cursed equilibrium, with an emphasis on how cursedness affects information revelation among players.

We begin with a discussion of the implications of cursed equilibrium in the context of voting. Recent rational-choice literature on voting in elections and juries assumes that people vote with a sophisticated understanding of how to predicate their votes on being pivotal. A voter who cares only about who wins an election should vote as if she is pivotal, even when she suspects that she is not.[25] Being pivotal can affect a voter's preferences if other voters have private information about the proper way to vote that is revealed from the fact that she is pivotal. Hence, a sophisticated voter asks herself what information other voters have that would make her pivotal, combines that information with her own private information, and votes accordingly.

In a series of papers, Feddersen and Pesendorfer (1996, 1997, 1998) explore the implications of such sophisticated reasoning by voters. Feddersen and Pesendorfer (1996) show that uninformed voters may strictly prefer abstaining to voting because they realize that if they are pivotal, they are more likely to decide the election in favor of the wrong candidate. By analogy to the winner's curse in auctions, they label this the "swing-voter's curse." The label is apt, since less than fully sophisticated voters may fall prey to such a curse much as bidders in common-values auctions fall prey to the winner's curse. To explore this possibility, cursed equilibrium was applied to the model developed in Feddersen and Pesendorfer (1998) of a jury deciding whether to convict a defendant of some crime. We discuss some general implications of cursed equilibrium in this model, as well as how our results fit the findings of Guarnaschelli, McKelvey, and Palfrey's (2000) experimental test of the model.

A jury of size $M \geq 2$ must decide whether to convict a defendant of a crime. Let ω_G be the state of the world where the defendant is guilty and let ω_I be the state of the world where the defendant is innocent, and suppose that jurors share the common prior $\mu(\omega_G) = \mu(\omega_I) = \frac{1}{2}$ Juror k receives a private signal $s_k \in \{\gamma, \iota\}$, correlated with the state of the world, with $\Pr[\gamma|\omega_G] = \Pr[\iota|\omega_I] = \theta \in (\frac{1}{2}, 1)$. Signals are independent conditional on the state of the world. Each juror k chooses an action $a_k \in \{g, i\}$, where g is a guilty vote and i is an innocent vote. Let $\sigma_k : \{\gamma, \iota\} \to \triangle\{g, i\}$ be k's strategy, which maps her signal to a probability distribution over guilty and innocent votes. Let n_G denote the number of jurors who vote guilty and let $n_I = M - n_G$ denote the number who vote innocent. Let $a \in \{A, C\}$ be the outcome of the jury process, where A denotes acquit and C denotes convict. The voting rule determines how the outcome depends on the jurors' votes. Under unanimous voting, the defendant is convicted if $n_G = M$; under majority voting, he is convicted if $n_G > n_I$. More generally, let $N \in [\frac{M}{2}, M]$ be the number of guilty votes needed to convict the defendant, so that the defendant is convicted if $n_G \geq N$. The role that different voting rules does (or does not) play in the likelihood of conviction is a focus of the literature in this area and is also a focus of our analysis.

[25]See Razin (2003) for a version of the sophisticated-voter model when voters care not just about who wins an election, but also about the margin of victory.

All jurors share the preferences

$$(14) \qquad u(a|\omega_G) = \begin{cases} q-1, & a = A, \\ 0, & a = C, \end{cases} \quad \text{and} \quad u(a|\omega_I) = \begin{cases} 0, & a = A, \\ -q, & a = C, \end{cases}$$

where $q \in (0, 1)$ is a parameter that measures the voters' trade-offs associated with either convicting the innocent or acquitting the guilty. The higher is q, the more jurors are bothered by convicting an innocent defendant relative to acquitting a guilty defendant. A juror prefers to convict if and only if she thinks the probability that the defendant is guilty exceeds q.

Given that the two states, ω_G and ω_I, are equally likely and that each private signal reflects the true state with probability $\theta > \frac{1}{2}$, a juror believes that the defendant is guilty with probability θ when her signal is γ and with probability $1-\theta$ when her signal is ι. We shall assume throughout that $1-\theta < q$, so that a juror who receives an innocent signal never votes to convict based on her information alone. In many applications, we shall consider the case of $q = \frac{1}{2}$, so an individual making a decision alone with only one signal would vote to convict if and only if the signal is guilty.

Because a juror's vote matters only when it is pivotal, it matters only if exactly $N-1$ other jurors cast guilty votes. The juror votes to convict if she thinks the probability of the defendant's being guilty is at least q given her own signal and the event that $N-1$ other jurors vote guilty. To find a symmetric equilibrium, consider the strategy σ_k, where

$$(15) \qquad \sigma_k(a_k = g|s_k) = \begin{cases} 1, & s_k = \gamma, \\ \sigma, & s_k = \iota, \end{cases}$$

for $\sigma \in [0, 1)$. Under strategy σ_k, Juror k votes guilty with probability 1 when she receives a guilty signal and votes guilty with probability σ when she receives an innocent signal. Feddersen and Pesendorfer show that a symmetric Bayesian Nash equilibrium of this form always exists. Of particular note is that the equilibrium often involves $\sigma > 0$, so that people with an innocent signal vote guilty with positive probability. To see the intuition for this result, note that if all those with innocent signals voted innocent, then a juror with one of those innocent signals should recognize that she is pivotal only when the defendant is almost certainly guilty. More generally, when q is low and $\frac{N}{M}$ is high, proper voting requires some of those with innocent signals to vote guilty.

It can be shown that the unique mixed-strategy χ-cursed equilibrium of the form described above is

$$(16) \qquad \sigma^* = \max\left\{0, \frac{\theta z - (1-\theta)}{\theta - (1-\theta)z}\right\}, \quad \text{where}$$

$$z = \left(\frac{1-q-\theta\chi}{q-(1-\theta)\chi}\right)^{1/(N-1)} \left(\frac{1-\theta}{\theta}\right)^{(M-N+1)/(N-1)}.$$

When $\theta > \frac{1}{2}$ and $q \geq \frac{1}{2}$, $\sigma^* > 0$ if

$$(17)\qquad \chi < \frac{1 - q - \left(\frac{1-\theta}{\theta}\right)^{2N-M+2}}{\theta - \left(\frac{1-\theta}{\theta}\right)^{2N-M+3}},$$

and $\sigma^* = 0$ otherwise.[26]

When χ is small and N is close to M, jurors with innocent signals vote guilty with positive probability, just as Feddersen and Pesendorfer found. More generally, cursed equilibrium shares many features of Bayesian Nash equilibrium. For each χ, $\frac{\partial\sigma^*}{\partial q} \leq 0$, meaning that the higher the burden of proof the jurors need to convict, the less likely they are to vote guilty. For all χ, $\frac{\partial\sigma^*}{\partial N} \geq 0$, meaning that the higher the number of guilty votes needed to convict, the more likely individual jurors are to vote guilty.

Although partially cursed jurors may vote strategically, they underinfer one another's information from pivotality. This affects their voting strategy and, hence, its efficiency—the likelihood that an innocent defendant is acquitted and a guilty defendant convicted. The formula above shows that $\frac{\partial\sigma^*}{\partial\chi} \leq 0$: because cursed jurors are less inclined to infer from the fact that they are pivotal that others have received guilty signals, cursedness causes jurors with innocent signals to be more likely to vote innocent. Indeed, when $\chi = 1$, voters simply vote their signals.

One striking result in Feddersen and Pesendorfer (1998) is that fixing the number of jurors, M, the probability of convicting an innocent defendant may increase in the number of guilty votes needed for conviction, N; this happens when the probability that a juror with an innocent signal votes guilty increases so much in response to a higher N that the odds of convicting an innocent defendant increase. Cursedness mitigates this connection. While increasing N can raise the probability of conviction even when $\chi > 0$, it decreases the probability of conviction for χ sufficiently close to 1, because in that case jurors with innocent signals always vote innocent irrespective of N.

While in the context of juries comparing unanimity rules to majority rules is natural, in large-scale elections it is of greater interest to compare intermediate cases where the share of votes needed to pass a proposition or elect a candidate is between $\frac{1}{2}$ and 1. Winning elections may typically require only a majority of votes, but passing a proposition often requires a supermajority such as two-thirds.[27] To consider the role of cursedness in such contexts, we consider the limit as M becomes very large and $N = kM$, where the fixed pa-

[26]Eyster and Rabin (2002) derive these and other results in more detail than we present here.

[27]In a multicandidate race with only two viable candidates, requiring a majority to avoid a runoff amounts de facto to requiring a supermajority.

rameter $k > \frac{1}{2}$ represents the percentage of guilty votes needed to "convict." In this case, it can be shown that

$$\lim_{\substack{M\to\infty \\ N=kM}} \sigma^* = \begin{cases} \dfrac{\left(\dfrac{1-\theta}{\theta}\right)^{(1-k)/k} - \dfrac{1-\theta}{\theta}}{1-\left(\dfrac{1-\theta}{\theta}\right)^{1/k}}, & \text{for } \chi < \dfrac{1-q}{\theta}, \\ 0, & \text{for } \chi > \dfrac{1-q}{\theta}. \end{cases} \tag{18}$$

When $\chi < \frac{1-q}{\theta}$, neither χ nor q affects the equilibrium proportion of guilty votes in the limit, but both χ and q help determine whether there is a mixed-strategy equilibrium in which voters with innocent signals sometimes vote guilty. Indeed, in the limit for $k < 1$, the election is fully efficient—always acquitting the innocent and convicting the guilty—if and only if the above mixed-strategy equilibrium exists. If the defendant is guilty, proportion $\theta + (1-\theta)\sigma^*$ of voters vote guilty, and if the defendant is innocent, proportion $(1-\theta) + \theta\sigma^*$ vote guilty. Voting is efficient when $(1-\theta) + \theta\sigma^* < k < \theta + (1-\theta)\sigma^*$. This holds for all values of $\theta > \frac{1}{2}$ and $k < 1$ when $\chi < \frac{1-q}{\theta}$. Note that $1-\theta < k < \theta$ holds even when $\sigma^* = 0$ if $\theta > k$. That is, if a higher percentage of voters get guilty signals than are needed to convict, guilty votes by those with innocent signals are not needed.

Given that whether $\sigma^* > 0$ is the sole determinant when $k > \theta$ of whether voting in large elections will be efficient, it is of special note that the condition for σ^* depends on χ but *not* on k. Since $\chi = 0$ always guarantees that $\sigma^* > 0$ when $k > \theta$, this means that any threshold election rule is efficient for large elections when voters are sufficiently uncursed. When $\chi > \frac{1-q}{\theta}$, by contrast, the election rule is efficient if and only if $\theta > k > \frac{1}{2}$; that is, the only election rules that guarantee efficiency for sufficiently cursed voters require conviction when voters vote naively.

A general principle is that voting mechanisms matter more for cursed than uncursed voters. Uncursed voters vote in a sophisticated manner by adjusting their behavior to whatever mechanism they face to assure as best they can that voting is efficient. By contrast, very cursed voters who vote based on their private information alone do not adjust their behavior to the mechanism to achieve efficiency. An efficient mechanism with cursed voters, therefore, needs to implement the right choice when voters vote naively. This suggests, in turn, that an efficient voting mechanism exists whenever there is a sufficiently large number of voters whose "naive preferences" depend on their private signals, so that aggregate voting behavior depends on whether the true state is that the defendant is guilty or innocent.

The only experimental test of the Feddersen and Pesendorfer model of which we are aware is Guarnaschelli, McKelvey, and Palfrey (2000), who study the laboratory behavior of randomly matched students at Caltech. Subjects were assigned to groups with either three or six members. Each group was assigned with equal probability to one of two urns, the "innocent" urn with seven innocent balls and three guilty balls, or the "guilty" urn with three innocent balls and seven guilty balls.[28] Subjects did not know to which urn their group had been assigned, but each subject privately and independently (sequentially with replacement) drew a ball at random from her group's urn. After observing her ball, each subject voted either innocent or guilty. Different groups had different rules for how their votes were aggregated. Subjects received 50 cents if their group's decision matched their urn and 5 cents if it did not. Guarnaschelli, McKelvey, and Palfrey's (2000) experiment corresponds to parameter values of $\mu(\omega_G) = \mu(\omega_I) = 0.5$, $q = 0.5$, and $\theta = 0.7$ in the model outlined above. They study the voting behavior in four different conditions: unanimous and majority rules in three- and six-person juries, i.e., $(N, M) \in \{(2, 3), (3, 3), (4, 6), (6, 6)\}$.

Subjects faced eight situations: four possible voting rules times two possible signals.[29] In six of the eight contingencies—in all cases where the observed signal is γ and the two majority-rule cases where the signal is ι—predicted behavior does not depend on χ. The first two lines of Table IV describe the two cases where it does: three- and six-person unanimous juries. Columns 4 and 5 give the percentage voting guilty in the Bayesian Nash and cursed equilibria, and column 6 shows the percentage of subjects actually voting guilty. As can be seen, too many people vote guilty in the three-person anonymous case—the *opposite* of the error that would be predicted by cursedness. On the other hand, too few people vote guilty in the six-person unanimous case, as predicted by cursedness. Column 7 indicates how each individual subject should vote if she knew how the others were voting. Column 8 shows the expected cost of the error in each case, in terms of lowering the expected likelihood of reaching the correct verdict from voting the wrong way, showing that the expected

TABLE IV

JURY VOTING[a]

M	N	s	$\sigma^*(0)$	$\sigma^*(1)$	$\widehat{\sigma}$	σ^{**}	Exp. Cost per Error	% Errors
3	3	ι	0.31	0.00	0.36	0.00	0.02	36
6	6	ι	0.65	0.00	0.48	1.00	0.03	52
Majority/ι			0.00	0.00	0.14	0.00	0.14	14
All/γ			1.00	1.00	0.95	1.00	0.20	5

[a]From Guarnaschelli, McKelvey, and Palfrey (2000).

[28]We follow the authors in using the language of "guilty" and "innocent," although the actual states described to the subjects were the more neutral "red" and "blue."

[29]The number of votes taking place in each of this eight situations varied between 143 and 202. In the two rows of Table III where we average across conditions, we take the simple average of the conditions, rather than weight by the number of subjects.

cost of those voting innocent in the six-person case is greater than the expected cost of those voting guilty in the three-person case.[30] Seen this way, column 9 shows that a higher percentage of subjects make this more costly error consistent with cursedness (52%) than make a less costly error inconsistent with cursedness (36%). This pattern is inconsistent with McKelvey and Palfrey's (1995) quantal-response equilibrium, where each player departs from best responding to others' behavior in such a way that more costly deviations occur less frequently than less costly ones.

As in Section 4, to estimate χ from the data, we enrich our basic solution concept by allowing players to make errors uncorrelated with cursedness: given how other subjects actually play the game, what value of χ makes players with innocent signals indifferent between voting innocent and guilty? As Table IV makes clear, cursedness provides no insight into the behavior of three-member juries, who vote guilty too frequently on innocent signals and too infrequently on guilty signals. Hence, a juror who best responds to others' behavior would never vote guilty on an innocent signal. Since cursedness increases only jurors' proclivity to vote their private information, they should vote innocent whatever χ is. Formally, $\chi = -0.17$ makes jurors with innocent signals indifferent between voting guilty and innocent. By contrast, in juries with six members, jurors vote guilty based on innocent signals too infrequently, as predicted by cursedness. Here, $\chi = 0.32$ makes jurors with innocent signals indifferent over guilty and innocent votes. Overall, in large juries subjects appear cursed, voting innocent on innocent signals more frequently than Bayesian Nash equilibrium suggests. In small juries, subjects play closer to Bayesian Nash equilibrium and even appear slightly too likely to convict with innocent signals. More evidence is needed on whether this result is robust, but these data alone provide little support for cursedness.

We now turn to apply cursed equilibrium to classical simple signalling games. Because it causes the receiver to infer less from signals than she should, a natural conjecture is that cursedness may make a high-quality type of sender unable to separate herself from a low-quality type by sending a costly signal and, hence, she is unwilling to send the signal. This intuition is not, however, always valid: because a cursed receiver does not fully infer that a sender who does not send a costly signal is a low type, cursedness may make a low type of sender less desperate to mimic a high type and, hence, make the high type able and willing to reveal herself by sending a costly signal.

To illustrate this, consider a situation where a sender is with equal probability one of two types, $t = b$ (bad) and $t = g$ (good). After learning her type, the sender can send one of two signals, $e = l$ (low) and $e = h$ (high). A receiver infers the sender's type from her signal, where $\widetilde{p}_l$ and $\widetilde{p}_h$ represent the receiver's beliefs about the probability that the sender is type g following signals l and h.

[30]If instead we compare the expected cost of the error conditional on being pivotal, the difference would be more dramatic: 19% versus 6% rather than 3% versus 2%.

After observing the signal the receiver chooses an action $a \in [0, 1]$ and has utility functions $u(a, g) = -(1-a)^2$ and $u(a, b) = -a^2$. The action a can be thought of as an investment that the receiver finds attractive if the sender is a good type but unattractive if he is a bad type. Hence, a receiver with beliefs $\widetilde{p}$ about the sender's type maximizes his expected utility $-\widetilde{p}(1-a)^2 - (1-\widetilde{p})a^2$ by choosing $a = \widetilde{p}_l$ and $a = \widetilde{p}_h$ following signals $e = l$ and $e = h$.

We assume that there is a continuous, increasing function $f : \mathbb{R} \to \mathbb{R}$ and real numbers $c_b > c_g > 0$ such that (presented in a reduced form that integrates the receiver's optimal response of $a = \widetilde{p}$) $u_b = u_g = f(\widetilde{p}_l)$ is the payoff to both types of sender if the signal l is sent, while $u_b = f(\widetilde{p}_h) - c_b$ and $u_g = f(\widetilde{p}_h) - c_g$ are the payoffs to the bad and good types of sender, respectively, if the signal h is sent. Thus, both types of sender want the receiver to believe that she is the good type; the signal h can potentially serve as a signal because it is more costly for the bad type than for the good type.

Because $c_b > c_g$, any separating Bayesian Nash equilibrium must involve type g sending signal h and type b sending l. For a separating equilibrium to exist, the good type must prefer to send h, so that $f(1) - c_g \geq f(0)$, and the bad type must prefer to send l, so that $f(1) - c_b \leq f(0)$. Hence, a separating Bayesian Nash equilibrium exists if and only if $c_g \leq f(1) - f(0) \leq c_b$.

When is there a separating χ-cursed equilibrium? In a separating equilibrium, because a χ-cursed receiver believes that type g sends h with probability $1 - \frac{\chi}{2}$ and type b sends h with probability $\frac{\chi}{2}$, he forms the beliefs $\widetilde{p}_l = \frac{\chi}{2}$ and $\widetilde{p}_h = 1 - \frac{\chi}{2}$. Hence, a separating χ-cursed equilibrium exists if and only if $c_g \leq f(1 - \frac{\chi}{2}) - f(\frac{\chi}{2}) \leq c_b$. When $\chi = 1$, $f(1 - \frac{\chi}{2}) - f(\frac{\chi}{2}) = f(\frac{1}{2}) - f(\frac{1}{2}) = 0$, so that no signalling can occur when the receiver is fully cursed. Intuitively, no sender would send a costly signal that would not affect the receiver's beliefs.[31]

While fully cursed receivers always destroy the potential for signalling, however, less extreme cursedness can *create* the potential for successful signalling. Indeed, if $c_b < f(1) - f(0)$, so that no separation can occur in a Bayesian Nash equilibrium, then because $f(1 - \frac{\chi}{2}) - f(\frac{\chi}{2})$ is decreasing in χ, there is some $\chi \in (0, 1)$ such that there is a separating cursed equilibrium. Intuitively, if the cost of being identified as the bad type is so high that the bad type prefers sending the costly signal to being identified, then full separation is not compatible with Bayesian Nash equilibrium. If the receiver is cursed enough that the bad type is just barely willing to behave differently than the good type, then the good type will be willing to reveal herself.[32]

We conclude this section with a theoretical and empirical analysis of an example of signalling that we call the revelation game, modelled after politicians

[31]While a separating Bayesian Nash equilibrium may not be a separating cursed equilibrium, recall that Proposition 3 demonstrates that every pooling Bayesian Nash equilibrium is a pooling χ-cursed equilibrium for every value of χ.

[32]Some similar implications play out in cheap-talk signalling games, as subsequently noted in Jehiel and Koessler (2005), who apply Jehiel's (2005) "analogy-based expectations equilibrium" to Bayesian games; a special case of their model coincides with fully cursed equilibrium.

who feel constrained not to lie to voters, but who do not feel constrained to reveal the full truth. In the 2000 American presidential campaign, President George W. Bush stated that he had never had an extramarital affair and had not used cocaine in the past 25 years, but he refused to say whether he used cocaine more than 25 years ago. Especially since President Bush volunteered the precise number 25 and his marital fidelity, fully rational voters probably should infer that Governor Bush used cocaine 26 years ago, but what would cursed voters infer from his (non)report?

Suppose a sender is of some type $t \in [0, 1]$, where t is a measure of her age the last time she engaged in some unseemly activity. A receiver does not know t, but has uniform priors on $[0, 1]$. A sender of type t chooses a message $m \in \{t, S\}$: she either announces her type or chooses S, meaning she remains silent.[33] After observing the sender's message, the receiver forms beliefs about the sender's type. We assume that the receiver picks an action $a(m) \in [0, 1]$ to maximize the expectation of his payoff, $-(a(m) - t)^2$. In equilibrium, the receiver chooses the action that coincides with his expectation of the sender's type given her message. The type t of sender's payoffs are $-a(m)$ if $m \in \{t, S\}$: she wants the receiver's beliefs to be as low as possible.

The unique perfect Bayesian equilibrium in this game is that all types reveal themselves fully. What are the cursed equilibria? Suppose the sender follows the cutoff strategy $r \in [0, 1]$, revealing her type iff $t < r$. A χ-cursed receiver forms beliefs $\chi\frac{1}{2} + (1 - \chi)(\frac{1}{2} + \frac{r}{2}) = \frac{1}{2} + (1 - \chi)\frac{r}{2}$, so the sender prefers to reveal whenever $t < \frac{1}{2} + (1 - \chi)\frac{r}{2}$. Because the marginal type r must be indifferent between revealing and not revealing, $r = \frac{1}{2} + (1 - \chi)\frac{r}{2}$, which implies $r = \frac{1}{1+\chi}$. Such a cutoff strategy is optimal for the sender, since types $t < r$ prefer revealing, while types $t > r$ prefer pooling.

When $\chi = 0$, $r = 1$ and all types reveal. The intuition is familiar: the lowest type always prefers to reveal herself. If only the lowest types reveal, then the lowest types who are supposed to pool will also prefer revealing, since they will have types lower than the average of all pooling types. For $\chi > 0$, however, some types pool. Because the receiver mistakenly believes that some types of sender who reveal pool and that some types of sender who pool reveal, when the receiver sees a sender who refuses to reveal her type, he thinks that she has a lower type, than she actually does.

An experiment by Forsythe, Isaac, and Palfrey (1989) provides evidence for cursedness in a version of the revelation game. In their experiment, each of four sellers was endowed with one unit of an object whose common value (in cents) to each of four bidders was drawn from a uniform distribution

[33]This game lies outside the purview of our formal definition of cursed equilibrium since players' action spaces depend on their types. Here we assume that when the cursed receiver receives the message $m = t$, she assigns probability 1 to the event that the sender's type is t. What matters for our qualitative results is that the cursed receiver puts lower weight on the average type when receiving the message t than S.

on $\{1, 2, \ldots, 125\}$. The sellers knew the value of their objects, but the bidders did not. The sellers chose whether to reveal the value of their objects to the bidders or conceal them; sellers could not misreport their values. Following this, the objects were auctioned to the bidders in first-price auctions. Just as in the revelation game, there is a cutoff χ-cursed equilibrium where sellers with objects valued more than $r = \frac{125\chi+1}{1+\chi}$ reveal their values and those with objects valued less than $r = \frac{125\chi+1}{1+\chi}$ conceal their values. Intuitively, low-value sellers conceal because cursed bidders overbid for objects with concealed values, mistakenly thinking that some high-value sellers conceal too. When $\chi = 0$, all sellers (except possibly those with the lowest possible valuation) reveal. When $\chi = 1$, sellers with valuations under 63 conceal and those with valuations above 63 reveal. Each bidder bids her expectation of the valuation of each seller's object, which is r for those sellers who conceal.

Forsythe, Isaac, and Palfrey ran 60 trials of this experiments with three groups of undergraduate subjects; the first group participated in 16 trials, and the second and third groups participated in 22 trials. Table V summarizes the data. The first row of the table shows the data for all sellers. For objects whose value was revealed, the winning bid was always approximately equal to the value of the object. Columns 2 and 3 show that 85 of 240 sellers (35%) concealed the value of their objects. The average value of concealed objects was 31, but the average winning bid was 39. Bidders suffered a significant winner's curse on blind-bid objects, paying more than their average value. The final row in the table shows that even subjects who had played the game more than ten times fell prey to the winner's curse.

We analyze all the data that allow for heterogeneity in subjects' χ and errors uncorrelated with cursedness in the same manner as in Section 4. We focus on bidder behavior, taking sellers' behavior as given.[34] A bidder i with parameter χ_i bids to maximize her payoff as if the expected value of the blind-bid

TABLE V

REVELATION GAME[a]

Group	Sellers	Conceal	Value Conceal	Bid Conceal
All	240	85	31	39
Experienced	120	32	23	27

[a] From Forsythe, Isaac, and Palfrey (1989).

[34] We could estimate a distribution of χ from seller behavior as well: to what distribution of χ are sellers best responding with error. If all bidders had the same χ, then a seller best responding to bidder behavior would blind bid iff her object were worth less than 39, meaning the average value of blind-bid objects would be 20. The fact that the actual average was 31 suggests that sellers overestimated the degree to which bidders were cursed.

object were her χ_i-virtual valuation

$$v_i(\chi_i) = (1-\chi_i)E[V|\text{blind bid}] + \chi_i E[V] = (1-\chi_i)38.8 + \chi_i 60.5. \tag{19}$$

Bidder's χ_i-virtual valuations are the χ_i-weighted average of the object's expected value unconditionally and conditional on being blind bid. Different bidders submit different bids for blind-bid objects for two reasons: different values of χ_i give them different χ_i-virtual valuations and the distributions of their opponents' bids differ. Let $F_{-i}(\cdot)$ be the distribution of the highest bid from bidders other than Bidder i and let $f_{-i}(b_i)$ be its density, estimated nonparametrically with a kernel estimator. Then Bidder i maximizes her expected payoff, $F_{-i}(b_i)(v_i(\chi_i) - b_i)$, at $\widehat{b_i} = v_i(\chi_i) - F_{-i}(\widehat{b_i})/f_{-i}(\widehat{b_i})$. Bidder i attempts to χ_i-cursedly best respond to the seller's empirical strategy and best responds to the empirical distribution of her opponents' bids. We assume that Bidder i bids her cursed best response plus error, $b_i = \widehat{b_i} + \varepsilon_i$, and estimate Bidder i's χ_i to be the value that minimizes the sum of squared errors across auctions.

Table VI reports our estimates as well as the same goodness-of-fit measure used in Section 4: assuming that all subjects have the same χ, we compute the sum of squared residuals between predicted and actual play. The first and second columns show that the data are bimodal. Half of subjects play close to the fully cursed best response to other subjects' play. The other half bid too high to be well described by our model. However, this pattern corresponds to the laboratory finding that even in private-values environments, subjects overbid in first-price auctions relative to Nash equilibrium (see Kagel (1995) for a

TABLE VI

ESTIMATED χ_i AND PREDICTION ERROR WITH UNITARY χ

χ	Observations	Squared Error (÷100,000)
−1.0	0	5.00
−0.8	0	4.93
−0.6	0	4.84
−0.4	0	4.82
−0.2	0	4.79
0.0	0	4.74
0.2	0	4.70
0.4	0	4.75
0.6	1	3.27
0.8	1	2.48
1.0	4	2.79
1.2	0	2.18
1.4	1	2.83
1.6	1	2.49
1.8	3	2.72
2.0	1	2.71

summary of the evidence). Column 3 shows that, for any value of $\chi \in (0, 1]$, our cursed-expectations concept fits the data better than $\chi = 0$, with a statistical tie between $\chi = 0$ and $\chi = 0.4$. As with the auction data in Section 4, any zero-degrees-of-freedom version of our model with any fixed $\chi > 0$ fits the data (weakly) better than Bayesian Nash equilibrium.

6. DISCUSSION AND CONCLUSION

We believe that cursed equilibrium can provide insight into many domains beyond those we analyze. For example, in organizational decision-making, cursedness may capture an exaggerated fear some parties have of delegation because they underappreciate how their delegates' future decisions depend on those delegates' information. Consider, for instance, a grand jury deciding whether to indict some defendant for a crime; an indictment moves the case to trial, where a jury hears the evidence and renders a verdict. A sufficiently cursed grand jury not fully convinced of the defendant's guilt may be too reluctant to indict because it fears the jury may convict an innocent defendant. Of course, it should realize that the jury convicts only if it has strong evidence that the defendant is guilty. Similar logic may play out in other organizations, where principals may be reluctant to delegate even to parties whose interests coincide with their own.

Some applications of cursed equilibrium point to its limitations and drawbacks, and we conclude by discussing some of these shortcomings and possible extensions. One limitation is that cursed equilibrium is defined only in games where each player's action space is independent of her type. In games without such independence, no one should believe that any type of any other player plays an action infeasible for that type. A problem with this approach, however, is that a cursed equilibrium in the game where an action is infeasible for a type of a player might differ from a cursed equilibrium in the related game where that same action is possible, but strictly dominated, for that type.

This problem, in turn, suggests modifying the definition to assume that no player thinks that any type of any other player plays a strictly dominated action in equilibrium. More generally, cursed equilibrium could be revised to incorporate the notion that the worse an action is for a type, the less likely other players think that type is to play it. Developing a new concept that incorporates this notion seems important both conceptually and for practical application, but would be inherently limited, since the very notion of cursed equilibrium is meant to capture limits to the degree to which people think through the relationship between others' relevant information and their behavior.

Another direction for development concerns a more important limitation to our current definition. Cursed equilibrium is meant to capture a general intuition that people underappreciate the relationship between other people's actions and private information. Yet the formal definition makes an artificial distinction between private information that is represented by type in a

Bayesian game and that which is not. In sequential games, for instance, our definition assumes that Player 3 does not fully appreciate how Player 2's actions depend on Player 2's types, but *does* fully appreciate how Player 2's actions depend on any actions that Player 2 observed Player 1 take but Player 3 did not. A more complete notion of cursed equilibrium would allow for "cursedness" over more general types of unobservable information. Treating "exogenous" and "endogenous" private information differently not only seems to us intuitively and psychologically wrong, but also creates some highly artificial differences in predictions based on the way a game is formally written down. In particular, insofar as a Bayesian game where one player has private information can be rewritten as another Bayesian game where a fictitious player takes an action observable only to the privately informed player, our definition of cursed equilibrium is not robust. Although we can think of no example where researchers have been or would be tempted to reframe games of interest with such fictitious players, the nonrobustness is conceptually troubling.

Another line of generalization would be to add more realistic variation in the degree of cursedness in different situations. For instance, players are probably more likely to ignore the informational content of other players' actions when they have not actually observed these actions than when they have; observing actions seems likely to induce more strategic sophistication. Hence, players in certain sequential games may be less cursed than they would be in corresponding simultaneous-move games. For example, Dekel and Piccione (2000) show in a rational model of binary voting that the set of informative equilibria does not depend on whether voters vote sequentially or simultaneously, which we conjecture also holds in cursed equilibrium. A psychologically richer approach to cursedness might incorporate the idea that a voter may better understand the relationship between other voters' signals and votes when she observes their votes than when she does not, leading to more rational voting in the sequential- than in the simultaneous-move voting procedure. Similarly, we believe that some experimental anomalies in sequential- versus simultaneous-move auctions (e.g., between English and second-price, sealed-bid auctions, or between Dutch and first-price, sealed-bid auctions) may arise from the salience of observing bids: people may have an easier time failing to appreciate the informational content of winning an auction if they must figure it out in the abstract than if they observe other bidders drop out before they win.

Dept. of Economics, London School of Economics, Houghton Street, London WC2A 2AE, U.K.; e.eyster@lse.ac.uk

and

Dept. of Economics, University of California–Berkeley, Berkeley, CA 94720, U.S.A.; rabin@econ.berkeley.edu; http://emlab.berkeley.edu/users/rabin/index.html.

Manuscript received November, 2000; final revision received February, 2005.

APPENDIX A: GENERALIZING THE MODEL

Cursed equilibrium can be generalized to (i) allow for nondegenerate distributions over different values of χ for the players, (ii) allow this distribution to differ for different players—and different types of each player, and (iii) allow for correlation in the distribution of cursedness across players and types of players. All three generalizations seem natural when players' cursedness depends on their experience. Different players may have different degrees of cursedness: in the no-trade setting, for instance, sellers of assets may be less cursed about any relevant private information than buyers, who on average may be less experienced. For the same reasons, different types of a player may have different distributions of cursedness: in the adverse-selection environment, for instance, buyers with higher valuations may have more experience because they have bought more such assets on average in the past. Finally, correlation in the players' cursedness follows naturally from correlation in their experience.

The generalization of the notion of cursed equilibrium to allow these asymmetries and heterogeneities is conceptually straightforward, although notationally cumbersome. In all the subsequent text we follow the notation of the text wherever meanings do not change or where the meaning of notation is obvious given the generalized context considered here. To begin, we assume a finite number of profiles of cursedness and let $f(\chi_1, \chi_2, \ldots, \chi_N; t_1, t_2, \ldots, t_N)$ be a probability mass function that gives the probability of any cursedness and type profile for the players.

Player k's strategy, $\sigma_k(\cdot|\chi_k, t_k)$, gives her probability distribution of actions a_k as a function of her type, t_k, and cursedness, χ_k. Given players' cursed-contingent strategies, for each *combination* of type and cursedness of each player, we are interested in her perception of the average strategy of the other players, averaged over their types. For each $t_k \in T_k$ and χ_k such that $f(\chi_k, \chi_{-k}; t_k, t_{-k}) > 0$ for some χ_{-k} and t_{-k}, let

$$g(\chi_{-k}, t_{-k}|t_k, \chi_k) = \frac{f(\chi_k, \chi_{-k}; t_k, t_{-k})}{\sum_{t'_{-k}, \chi'_{-k}} f(\chi_k, \chi'_{-k}; t_k, t'_{-k})}$$

be the probability of (χ_{-k}, t_{-k}) conditional on Player k's type and cursedness, (t_k, χ_k). We define

$$\overline{\sigma}_{-k}(a_{-k}|t_k, \chi_k) \equiv \sum_{t_{-k}, \chi_{-k}} g(\chi_{-k}, t_{-k}|t_k, \chi_k) \cdot \sigma_{-k}(a_{-k}|t_{-k}, \chi_{-k}). \tag{20}$$

When Player k is of type t_k and cursedness χ_k, $\overline{\sigma}_{-k}(a_{-k}|t_k, \chi_k)$ is the probability that players $j \neq k$ play action profile a_{-k} when they follow strategy σ_{-k}. Its interpretation coincides with $\overline{\sigma}_{-k}(a_{-k}|t_k)$ in the text: a player who (mistakenly) believes that each type profile of the other players plays the same mixed action profile believes that the other players are playing $\overline{\sigma}_{-k}(\cdot|t_k, \chi_k)$ whenever they play $\sigma_{-k}(a_{-k}|t_{-k})$. Note that the dependence of $\overline{\sigma}_{-k}(\cdot|t_k,\chi_k)$ on χ_k does not indicate that the degree of misprediction implicit in $\overline{\sigma}_{-k}(\cdot|t_k,\chi_k)$ corresponds to χ_k; as in the text, we mean here that $\overline{\sigma}_{-k}(\cdot|t_k,\chi_k)$ would be Player k's beliefs about the strategies of other players if she were fully cursed. Rather, χ_k affects $\overline{\sigma}_{-k}(\cdot|t_k,\chi_k)$ only insofar as it may be correlated with χ_{-k}, which may affect the $-k$'s strategy.

DEFINITION 5: A mixed-strategy profile σ is an *f-generalized cursed equilibrium* if for each k, $t_k \in T_k$ and each χ_k in the support of f, and each a_k^* such that $\sigma_k(a_k^*|t_k,\chi_k) > 0$,

$$a_k^* \in \arg\max_{a_k \in A_k} \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k}\in A_{-k}} \big[\chi_k \overline{\sigma}_{-k}(a_{-k}|t_k,\chi_k) + (1-\chi_k)\sigma_{-k}(a_{-k}|t_{-k})\big] \times u_k(a_k, a_{-k}; t_k, t_{-k}).$$

All the natural generalizations of the results and proofs in Section 2 hold. It is worth noting that we make no implicit assumptions about players' awareness of other players' cursedness. In equilibrium—a very strong concept here, as always—an uncursed player would correctly predict other players' type-contingent strategies, regardless of her beliefs about their cursedness. In similar fashion, although a cursed player misinterprets other players' type-contingent strategies, she does it in a way that requires no assumptions about her beliefs about their cursedness.

APPENDIX B: PROOFS

PROOF OF LEMMA 1: From Bayes' rule,

$$\begin{aligned}(21)\qquad &\widehat{p}_{t_k}(t_{-k}|a_{-k},\sigma_{-k})\\ &= \frac{(1-\chi)\sigma_{-k}(a_{-k}|t_{-k}) + \chi\overline{\sigma}_{-k}(a_{-k}|t_k)}{\sum_{t'_{-k}\in T_{-k}}((1-\chi)\sigma_{-k}(a_{-k}|t'_{-k}) + \chi\overline{\sigma}_{-k}(a_{-k}|t_k))p(t'_{-k}|t_k)}\\ &\quad\times p(t_{-k}|t_k)\\ &= \left((1-\chi)\frac{\sigma_{-k}(a_{-k}|t_{-k})}{\overline{\sigma}_{-k}(a_{-k}|t_k)} + \chi\right)p(t_{-k}|t_k).\end{aligned}$$

Q.E.D.

PROOF OF PROPOSITION 1: Consider the alternative game $\overline{G}^{\chi} \equiv (A, T, p, \overline{u}^{\chi})$, where (A, T, p) are all the same, but u is replaced by

$$(22)\qquad \overline{u}_k^{\chi}(a_k, a_{-k}; t_k, t_{-k}) \equiv (1-\chi)u_k(a_k, a_{-k}; t_k, t_{-k}) + \chi \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)u_k(a_k, a_{-k}; t_k, t_{-k}).$$

The utility function of type t_k of Player k is the χ-weighted average of her actual utility function and her "average utility function," averaged over all possible types of her opponents. The σ denotes a Bayesian Nash equilibrium of $\overline{G}^\chi$ if for each Player k, each type $t_k \in T_k$, and each a_k^* such that $\sigma_k(a_k^*|t_k) > 0$,

$$
\begin{aligned}
(23)\qquad a_k^* &\in \arg\max_{a_k \in A_k} \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k}\in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) \overline{u}_k^\chi(a_k, a_{-k}; t_k, t_{-k}) \\
&= (1-\chi) \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k}\in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) u_k(a_k, a_{-k}; t_k, t_{-k}) \\
&\quad + \chi \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k}\in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) \\
&\qquad \times \sum_{t'_{-k}\in T_{-k}} p_k(t'_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t'_{-k}).
\end{aligned}
$$

However,

$$
\begin{aligned}
(24)\qquad & \chi \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k}\in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) \\
&\qquad \times \sum_{t'_{-k}\in T_{-k}} p_k(t'_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t'_{-k}) \\
&= \chi \sum_{a_{-k}\in A_{-k}} \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sigma_{-k}(a_{-k}|t_{-k}) \\
&\qquad \times \sum_{t'_{-k}\in T_{-k}} p_k(t'_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t'_{-k}) \\
&= \chi \sum_{a_{-k}\in A_{-k}} \sum_{t'_{-k}\in T_{-k}} p_k(t'_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}) \overline{\sigma}_{-k}(a_{-k}|t_k) \\
&= \chi \sum_{t'_{-k}\in T_{-k}} p_k(t'_{-k}|t_k) \sum_{a_{-k}\in A_{-k}} \overline{\sigma}_{-k}(a_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}),
\end{aligned}
$$

and, hence,

$$
\begin{aligned}
(25)\qquad & \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k}\in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) \overline{u}_k^\chi(a_k, a_{-k}; t_k, t_{-k}) \\
&= \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k}\in A_k} \left[\chi \overline{\sigma}_{-k}(a_{-k}|t_k) + (1-\chi)\sigma_{-k}(a_{-k}|t_{-k})\right] \\
&\qquad \times u_k(a_k, a_{-k}; t_k, t_{-k}).
\end{aligned}
$$

Thus if σ is a Bayesian Nash equilibrium of $\overline{G}^{\chi}$, it is also a cursed equilibrium of G. Because $\overline{G}^{\chi}$ is finite, it has a Bayesian Nash equilibrium, and so G has a cursed equilibrium. *Q.E.D.*

PROOF OF PROPOSITION 2: If each type t_k of each player k's expected payoff from playing a_k when the other players play a_{-k} in the χ-virtual game $\overline{G}^{\chi}$ is independent of χ, then the result follows since the set of Bayesian Nash equilibria of $\overline{G}^0 = G$ coincides with the set of Bayesian Nash equilibria of $\overline{G}^{\chi}$, which by Proposition 1 is the set of χ-cursed equilibria of G. Hence it suffices to show that

$$\begin{aligned}(26)\qquad &\sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)u_k(a_k, a_{-k}; t_k, t_{-k})\\ &= \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)u_k(a_k, a_{-k}; t_k, t_{-k}).\end{aligned}$$

The second expression can be rewritten

$$\begin{aligned}(27)\qquad &\sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{t_{-0k}\in T_{-0k}} p_k(t_{-0k}|t_k)\\ &\qquad \times \sum_{t_0\in T_0} p_k(t_0|t_k, t_{-0k})u_k(a_k, a_{-k}; t_0, t_k, t_{-0k})\\ &= \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{t_0\in T_0} p_k(t_0|t_k, t_{-0k})u_k(a_k, a_{-k}; t_0, t_k, t_{-0k})\\ &\qquad \times \sum_{t_{-0k}\in T_{-0k}} p_k(t_{-0k}|t_k),\end{aligned}$$

since $E[u_k(a_k, a_{-k}; t_0, t_k, t_{-0k})|t_k, t_{-0k}]$ is independent of t_{-0k}. It simplifies to

$$\begin{aligned}(28)\qquad &= \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \sum_{t_0\in T_0} p_k(t_0|t_k, t_{-0k})u_k(a_k, a_{-k}; t_0, t_k, t_{-0k})\\ &= \sum_{t_{-0k}\in T_{-0k}} p_k(t_{-0k}|t_k) \sum_{t_0\in T_0} p_k(t_0|t_k, t_{-0k})u_k(a_k, a_{-k}; t_0, t_k, t_{-0k})\\ &= \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)u_k(a_k, a_{-k}; t_k, t_{-k})\end{aligned}$$

as desired. *Q.E.D.*

PROOF OF PROPOSITION 3: Suppose that σ is a strategy profile such that for each Player k there exists some $a_k \in A_k$ such for each $t_k \in T_k$ $\sigma(a_k|t_k) = 1$.

Then

$$\overline{\sigma}_{-k}(a_{-k}|t_k) \equiv \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)\sigma_{-k}(a_{-k}|t_{-k}) = \sigma_{-k}(a_{-k}|t_{-k}) \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) = \sigma_{-k}(a_{-k}|t_{-k}), \tag{29}$$

since $\sigma_{-k}(a_{-k}|t_{-k})$ does not depend on t_{-k}. If σ is a χ-cursed equilibrium, then a_k maximizes

$$\sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \cdot \sum_{a_{-k}\in A_{-k}} \big[\chi\overline{\sigma}_{-k}(a_{-k}|t_k) + (1-\chi)\sigma_{-k}(a_{-k}|t_{-k})\big] \times u_k(a_k, a_{-k}; t_k, t_{-k}) = \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \cdot \sum_{a_{-k}\in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k})u_k(a_k, a_{-k}; t_k, t_{-k}), \tag{30}$$

which does not depend on χ. Therefore, whatever χ, a_k maximizes Player k's expected payoff given that players $j \neq k$ play $\sigma_{-k}(a_{-k}|t_{-k})$, so σ is a χ-cursed equilibrium for every $\chi \in [0, 1]$. *Q.E.D.*

The proofs of Propositions 4 and 5 were outlined in the text.

PROOF OF PROPOSITION 6: Proposition 5 establishes that the symmetric χ-cursed equilibrium of the second-price auction has $b_i(x_i) = (1-\chi)v^N(x_i, x_i) + \chi r^N(x_i)$. It suffices to show that the seller's expected revenue is

$$(1-\chi)E\big[v^N\big(Y^N(2), Y^N(2)\big)\big] + \chi E\big[r^N(Y^N(2))\big] > E[u^N(X^N, S)], \tag{31}$$

where $Y^N(2)$ is the second-highest of the N bidders' signals. Let $Y_s^N(2)$ be the second-highest signal conditional on the event that $S = s$. For each s, $Y_s^N(2) \to \overline{x}(s)$ in probability. Since $u^N(x, s) \geq 0$ for each x and s is continuous, given any $\varepsilon > 0$, for sufficiently large N,

$$E\big[r^N(Y^N(2))\big] = \int r^N(Y_s^N)h(s)\,ds > E\big[r^N(\overline{x}(S))\big] + \frac{\varepsilon}{2}.$$

By assumption, there exists some $\varepsilon > 0$ such that for each N, $E[r^N(\overline{x}(S))] > E[u^N(X, S)] + \frac{\varepsilon}{2}$, which implies that for sufficiently large N, $E[r^N(Y^N(2))] - E[u^N(X^N, S)] > \varepsilon$. Because the sequence (A^N) is competitive, for large N,

$$E\big[v^N\big(Y^N(2), Y^N(2)\big)\big] - E[u^N(X^N, S)] > -\frac{\chi\varepsilon}{(1-\chi)}. \tag{32}$$

Hence,

$$(33)\qquad \chi\big(E\big[r^N(Y^N(2))\big] - E[u^N(X,S)]\big) + (1-\chi)\big(E\big[v^N\big(Y^N(2), Y^N(2)\big)\big] - E[u^N(X^N,S)]\big) > 0,$$

which establishes the result. *Q.E.D.*

PROOF OF PROPOSITION 7: Since r is symmetric, it is straightforward to show that $E[r(Y^N(2))] \geq E[r(X)] = E[u(X,S)] \geq E[v(Y^N(2), Y^N(2))]$ (where the first inequality holds with equality for $N = 3$). Since the seller's expected revenue is a χ-weighted average of the first and last terms, it increases in χ. *Q.E.D.*

PROOF OF PROPOSITION 8: For $N \leq 3$ and r symmetric, $E[r(Y^N(2))] \leq E[r(X)] = E[u(X,S)]$. Since $E[v(Y^N(2), Y^N(2))] \leq E[u(X,S)]$, the seller's expected revenue is a χ-weighted average of two terms less than the average value of the object. *Q.E.D.*

REFERENCES

AKERLOF, G. (1970): "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84, 488–500.

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.

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

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.

DEKEL, E., D. FUDENBERG, AND D. K. LEVINE (2004): "Learning to Play Bayesian Games," *Games and Economic Behavior*, 46, 282–303.

DEKEL, E., AND M. PICCIONE (2000): "Sequential Voting Procedures in Symmetric Binary Elections," *Journal of Political Economy*, 108, 34–55.

EYSTER, E., AND M. RABIN (2002): "Cursed Equilibrium," Working Paper E02-320, IBER, UC Berkeley.

FEDDERSEN, T., AND W. PESENDORFER (1996): "The Swing Voter's Curse," *American Economic Review*, 86, 408–424.

——— (1997): "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica*, 65, 1029–1058.

——— (1998): "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting," *American Political Science Review*, 90, 23–35.

FORSYTHE, R., R. M. ISAAC, AND T. R. PALFREY (1989): "Theories and Tests of 'Blind Bidding' in Sealed-Bid Auctions," *RAND Journal of Economics*, 20, 214–238.

GUARNASCHELLI, S., R. MCKELVEY, AND T. PALFREY (2000): "An Experimental Study of Jury Decision Rules," *American Political Science Review*, 94, 407–423.

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

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

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

JACOBSEN, B., J. POTTERS, A. SCHRAM, F. VAN WINDEN, AND J. WIT (2000): "(In)accuracy of a European Political Stock Market: The Influence of Common Value Structures," *European Economic Review*, 44, 205–230.

JEHIEL, P. (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, forthcoming.

JEHIEL, P., AND F. KOESSLER (2005): "Revisiting Games of Incomplete Information with Analogy-Based Expectations," Unpublished Manuscript.

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.

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

——— (2002): *Common Values Auctions and the Winner's Curse*. Princeton, NJ: Princeton University Press.

LEVIN, D., J. H. KAGEL, AND J.-F. RICHARD (1996): "Revenue Effects and Processing in English Common Value Auctions," *American Economic Review*, 86, 442–460.

MCKELVEY, R. D., AND T. R. PALFREY (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 7, 6–38.

MILGROM, P., AND N. STOKEY (1982): "Information, Trade, and Common Knowledge," *Journal of Economic Theory*, 26, 17–27.

MILGROM, P. R., AND R. J. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

POTTERS, J., AND J. WIT (1995): "Bets and Bids: Favorite-Longshot Bias and Winner's Curse," Discussion Paper 9604, CentER, Tilburg University.

RAZIN, R. (2003): "Signalling and Electing Motivations in a Voting Model with Common Values and Responsive Candidates," *Econometrica*, 71, 1083–1120.

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.

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

VON GOETHE, J. W. (1808, 1832; 1949, 1959): *Faust*, Part I, transl. by P. Wayne (Vols. 1 and 2). Harmondsworth, U.K.: Penguin.

WILDE, O. (1891, 1982): *The Picture of Dorian Gray*, ed. by P. Ackroyd. Harmondsworth, U.K.: Penguin.