

University of California, Berkeley
Department of Economics
Working Paper No. E02-320

Cursed Equilibrium

Erik Eyster

Nuffield College, Oxford

Matt Rabin

University of California, Berkeley

August 2002 (Nov 2000)

Keywords: curses

JEL Classification: B49

Abstract

There is evidence that people do not fully take into account how other people's actions are contingent on these others' 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 these other players' information. We apply the concept to common-values auctions, where cursed equilibrium captures the widely observed phenomenon of the winner's curse. We also show how cursed equilibrium predicts other empirically observed phenomena, such as trade in adverse-selection settings where conventional analysis predicts no trade, and "naïve" voting in elections and juries where rational-choice models predict that voters fully take into account the informational content in being pivotal.

We thank Colin Camerer and seminar participants at Berkeley, Cal Tech, and Nuffield College (Oxford) for helpful comments, and Davis Beekman, Kitt Carpenter, and David Huffman for valuable research assistance. Eyster thanks the Olin and MacArthur Foundations, and Rabin thanks the Russell Sage, Sloan, MacArthur, and National Science Foundations for financial support. Both authors deny using cocaine less than 26 years before working on this paper.

Contacts: E-mail: erik.eyster@nuffield.oxford.ac.uk and rabin@econ.berkeley.edu. Rabin's web page: <http://elsa.berkeley.edu/rabin/index.html>.

This paper is available on-line at the California digital Library/ escholarship website: <http://repositories.cdlib.org/iber/econ/> and at the publications website for the University of California, Berkeley, Department of Economics: <http://iber.berkeley.edu/wps/econwp.html>

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 for this phenomenon 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’ private information 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 make this error.¹ Indeed, empirical researchers base their estimations of bidders’ valuations for the object being auctioned on the presumption that bidders do not make this error.² 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 which 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 (1988) and Holt and Sherman (1994) both by allowing players to partially, but not fully, appreciate the information content in other players’ actions, and by defining a solution concept applicable to general Bayesian games. We flesh out the implications of our model in common-values auctions and many other settings, discussing the empirical evidence that motivates it in the specific contexts we consider. The model ties together a wide range of empirically observed phenomena with a formalization of a single psychological principle — the underappreciation of the informational content of the behavior of others.

¹See Thaler (1988) for an overview of the early evidence on the winner’s curse as well as Kagel (1995) for a survey of laboratory auctions.

²In fact, when “winner’s curse” appears in the title of a paper, it typically refers to the study of players who avoid rather than succumb to the curse. Just as suburban housing developments are often named after the bit of nature obliterated to create them (“Forest Glen”), so too the term winner’s curse is typically used to describe what isn’t there.

³Potters 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.

In Section 2 we present our equilibrium concept. We consider standard Bayesian games where players' private information is represented by their "types," whose joint distribution is common knowledge. 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 average action of what all types of other players are playing, rather than their true, type-specific action. Players choose their actions to maximize their expected utilities 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 action rather than their type-contingent strategy. 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 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 comes in misunderstanding the relationship 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. The seller knows whether the car is a lemon, worth \$0 to both the seller and buyer, or a peach, worth \$3,000 to the buyer and \$2,000 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 will sell if and only if the car is a lemon, and hence refuse to buy, a cursed buyer may mistakenly buy the car. The sure sale of the lemon is a χ -cursed equilibrium because a χ -cursed buyer believes that with probability χ the seller sells irrespective of the type of car, so that the car being sold is a peach with probability with $(1 - \chi) \cdot 0 + \chi \cdot \frac{1}{2} = \frac{\chi}{2}$, and therefore worth $\frac{\chi}{2} \cdot 3,000 = 1,500\chi$. Hence, a buyer cursed by $\chi > \frac{2}{3}$ will buy the car, only to discover that whenever the seller is willing to sell it is a worthless lemon.

We prove that every finite game has (for every value of χ) a cursed equilibrium — by observing that a cursed equilibrium corresponds to a Bayesian Nash equilibrium in a modified game where the players' payoffs for each action and type profile are a weighted average of their actual payoffs and their average payoffs for that action profile averaged over other players' types. We also show that when each player's payoffs are fully independent of other players' types, cursed equilibrium and Bayesian Nash equilibrium coincide. Intuitively, the only difference between the two equilibrium 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, then such mistaken beliefs do not matter. Finally, we define a *perfectly-cursed equilibrium*, the analogue to Perfect Bayesian equilibrium, and show how it imposes an important restriction on players' beliefs off the equilibrium path.

In Sections 3, 4, and 5, we apply the general model to three different important settings — bilateral trade, auctions, and voting. Our model both helps to 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. But 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 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.

In Section 4 we turn to our primary motivating application, common-values auctions. In a cursed equilibrium, bidders in a symmetric equilibrium do not recognize that they win the auction only when they have the most positive information about the value of the object. When χ and the number of bidders are high enough, this leads to the winner's curse — the average winning bid exceeds the average value of the object. Even though cursed bidders may suffer the winner's curse, while rational bidders never do, we show that cursedness does not always raise the seller's expected revenue, because cursed bidders may also sometimes bid less than rational bidders. Finally, we compare the predictions of cursed equilibrium to some of the experimental evidence on common-values auctions.

In Section 5 we apply cursed equilibrium to a model of voting, contrasting our predictions to those of a recent rational-choice literature on voting in elections and on juries. This literature

assumes that people vote with a sophisticated understanding that they should predicate their votes on being pivotal, which means a voter should vote not based on her beliefs at the time of voting, but rather based on what her beliefs would be if her vote decided the election. Just as in bidding, therefore, voters must predict the relationship between other voters’ private information and their votes. We show that because of this underinference, cursed voters are more likely to vote “naively” 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. We also discuss whether cursed equilibrium can help explain McKelvey and Palfrey’s (1998) findings in their experimental test of jury voting.

In Section 6, we illustrate the implications of cursed equilibrium in two different signaling contexts. First, we consider classical simple signaling games, where fully-cursed equilibrium rules out the use of costly signaling, but lesser degrees of cursedness can either destroy meaningful signaling arising in a Bayesian Nash equilibrium or facilitate meaningful signaling that could not arise in a Bayesian Nash equilibrium. Second, we apply cursed equilibrium to a model of “verifiable cheap talk” modeled 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, one Bayesian Nash equilibrium is for each type of politician to reveal her type, since any politician who knows the truth to be less damaging than fully rational voters infer from silence prefers to reveal. 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.

Because it posits that each player correctly predicts the equilibrium distribution of other players’ actions without correctly predicting their type-contingent strategies, cursed equilibrium is incompatible with many natural explanations for how equilibrium play arises. In some settings, however, we think this is a natural occurrence: a player who observes repeated play of a single game may learn the distribution of other players’ actions, but because she may never observe other players’ private information she may not learn the relationship between the other players’ actions and private information. While we do not find this foundation for our approach fully satisfying, we believe cursed equilibrium provides a useful, general, and relatively non-arbitrary way to study the behavioral implications of a pervasive form of failure of contingent thinking.

Our formulation of cursed equilibrium is an over-simplification in many other ways that may limit its applicability beyond the set of games we consider. In some contexts, our formulation may make some unrealistic predictions; in others, cursed equilibrium is not well-defined. We conclude the paper in Section 7 with a discussion of possible extensions of the notion of cursed equilibrium that might cope with these problems, as well as discussing some possible further economic applications of the principles developed in this paper.

2 Definition and General Results

Before developing specific applications, in this section we formally define cursed equilibrium, prove its existence in all finite Bayesian games, and develop some general principles and results. Consider a finite Bayesian Game, $G = (A_1, \dots, A_N; T_0, T_1, \dots, T_N; p; u_1, \dots, u_N)$, played by players $k \in \{1, \dots, N\}$. A_k is 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 in our analysis below, we include a set of "nature's types", T_0 . $T \equiv T_0 \times T_1 \times \dots \times T_N$ is the set of type profiles, and p is the probability distribution over T , which we assume is common to all players. Player k 's payoff function $u_k : A \times T \rightarrow \mathbb{R}$ depends on all players' actions $A \equiv A_1 \times \dots \times A_N$ and their types. A (mixed) strategy σ_k for Player k specifies a probability distribution over actions for each type: $\sigma_k : T_k \rightarrow \Delta A_k$. Let $\sigma_k(a_k|t_k)$ be the probability that type t_k plays action a_k , and let $u \equiv (u_1, \dots, u_N)$.

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 $T_{-k} \equiv \prod_{j \neq k} T_j$ of other players (including nature) given her own type $t_k \in T_k$. Let $A_{-k} \equiv \prod_{j \neq 0, k} A_j$ be the set of action profiles for players $j \neq k$ (excluding nature, who takes no action), and $\sigma_{-k} : T_{-k} \rightarrow \prod_{j \neq 0, k} \Delta A_j$ 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 $\sigma_k(a_k^*|t_k) > 0$,*

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

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.

Before 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 $\bar{\sigma}_{-k}(\cdot|t_k)$ by

$$\bar{\sigma}_{-k}(a_{-k}|t_k) \equiv \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \cdot \sigma_{-k}(a_{-k}|t_{-k}).$$

When Player k is of type t_k , $\bar{\sigma}_{-k}(a_{-k}|t_k)$ 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 $\bar{\sigma}_{-k}(\cdot|t_k)$ whenever they play $\sigma_{-k}(a_{-k}|t_{-k})$. Note that $\bar{\sigma}_{-k}(a_{-k}|t_k)$ depends on t_k , so different types of Player k have different beliefs about the average action of players $j \neq k$. Let $\bar{\sigma}_{-k}(t_k) : T_{-k} \rightarrow \prod_{j \neq 0, k} \Delta A_j$ denote t_k 's beliefs about the average strategy of players $j \neq k$, namely $\bar{\sigma}_{-k}(t_k)$ is the strategy players $j \neq k$ would play if each type profile t_{-k} played a_{-k} with probability $\bar{\sigma}_{-k}(a_{-k}|t_k)$.

From this, we define a cursed equilibrium, defined 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 $\sigma_k(a_k^*|t_k) > 0$,*

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

In a χ -cursed equilibrium, each player correctly predicts the probability distribution over her opponents' actions, but she misunderstands the relationship 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.⁴ 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

⁴To 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

$$\begin{aligned} & \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) [\chi \bar{\sigma}_{-k}(a_{-k}|t_k) + (1 - \chi) \sigma_{-k}(a_{-k}|t_{-k})] \\ = & \chi \bar{\sigma}_{-k}(a_{-k}|t_k) + (1 - \chi) \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \sigma_{-k}(a_{-k}|t_{-k}) = \bar{\sigma}_{-k}(a_{-k}|t_k). \end{aligned}$$

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*.

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.⁵ 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 Players j and k share common beliefs about Player l ’s strategy. In many of our applications, however, players’ types are not independent, so that differences in beliefs prevail in equilibrium.⁶

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 discussing auctions, for instance, we often 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.⁷

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

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

When $\chi = 0$, $\widehat{p}_{t_k}(t_{-k}|a_{-k}, \sigma_{-k}) = \frac{\sigma_{-k}(a_{-k}|t_{-k})}{\bar{\sigma}(a_{-k}|t_k)} p_k(t_{-k}|t_k)$: Player k correctly updates her beliefs about the other players according to Bayes Rule. When $\chi = 1$, $\widehat{p}_{t_k}(t_{-k}|a_{-k}, \sigma_{-k}) = p_k(t_{-k}|t_k)$:

⁵In 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. But, 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.

⁶For example, suppose that there are two possible states of the world, ω_1 and ω_2 , and each player receives one of two possible signals, s_1 and s_2 , where $\Pr[s_i|\omega_i] > \Pr[s_i|\omega_j]$. Suppose that in equilibrium each player takes action a_i if her signal is s_i . When she receives signal s_1 , Player 1 thinks ω_1 more likely than she did before receiving her signal, and therefore she thinks it more likely that Player 2 also receives signal s_1 . In both a Bayesian Nash equilibrium and a cursed equilibrium, a Player 1 with signal s_1 thinks it more likely that Player 2 takes action a_1 than a Player 1 with signal s_2 . In a χ -cursed equilibrium, however, a Player 1 with signal s_1 also thinks it more likely that Player 2 plays a_1 when his signal is s_2 because the average probability that Player 2 plays a_1 is higher when Player 1’s signal is s_1 than s_2 .

⁷All proofs are in the Appendix.

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 demonstrates that in finite games, where Bayesian Nash equilibria exist, χ -cursed equilibria also exist.

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 explicated 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

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

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 other 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 $\bar{G}^\chi \equiv (A, T, p, \bar{u}^\chi)$, G has a cursed equilibrium whenever \bar{G}^χ has a Bayesian Nash equilibrium. We use this reinterpretation and alternative formalization of cursed equilibrium as the Bayesian Nash equilibrium of \bar{G}^χ repeatedly below.

Proposition 1 follows from the fact that whenever G is finite, \bar{G}^χ is finite, and finite games have at least one Bayesian Nash equilibrium. Proposition 1 is of limited general interest, however. While every game we consider in this paper has an equilibrium for each value of χ , most of the games we consider have uncountably infinite type and action spaces, so Proposition 1 does not guarantee existence in these games. Moreover, the existence of a Bayesian Nash equilibrium ($\chi = 0$) is neither necessary nor sufficient for the existence of a χ -cursed equilibrium for each $\chi \in (0, 1]$. However, the counterexamples we have devised to show this involve games with discontinuous payoffs or non-

compact action spaces, and we suspect that in well-behaved games where Bayesian Nash equilibria exist cursed equilibria also exist.⁸

In a cursed equilibrium, a player maximizes her payoffs under the mistaken belief that other players' actions depend less on their types than they 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 need to distinguish between the set of Player k 's opponents and the set of possible states of the world. Let $T_{-0k} \equiv \prod_{i \neq 0, k} T_i$ 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 any other player's type, but also that no player can learn anything about her expected payoff by learning any other 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. But if one bidder learned another bidder's signal her beliefs about the value of the object, and therefore her beliefs about her payoffs from a profile of bids, would change. Hence, Proposition 3 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.

⁸ Athey (1997) proves an existence theorem for infinite games satisfying a single-crossing property: if each player's best response to every strategy of her opponents that is increasing in their types is increasing in her type (and payoffs are continuous in actions), then the game has a pure-strategy equilibrium where each player's strategy is increasing in her type. While space constraints prevent us from proving it in this paper, the same is true of cursed equilibrium: a game that satisfies Athey's conditions has an increasing, pure-strategy χ -cursed equilibrium for each value of χ . (Likewise, Milgrom and Roberts' (1990) monotone-comparative-statics results for Nash equilibria in supermodular games apply to cursed equilibria.)

The intuition behind the proposition is that if a player learns nothing about her expected payoff from knowing the other players' types, then it does not matter that she misunderstands the relationship between the other players' types and their actions. More precisely, a player who correctly predicts the probability distribution over the other players' actions who does not learn anything about her expected payoff from learning the other players' types chooses the same action irrespective of her theory of which types of the other players play which action.

A final result is of interest in some applications and helps give more intuition about the nature of the cursed equilibrium. By analogy with pooling equilibria in simple signaling 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 $t_k \in T_k$, $\sigma(a_k|t_k) = 1$. Then:

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. This is because in a pooling equilibrium players' actions are independent of their types. Ignoring the relationship between others' actions and their information is not a mistake when there *is* no relationship.

In many Bayesian games, especially sequential games, researchers apply refinements of Bayesian Nash equilibrium in making predictions. A simple way to define analogous refinements of χ -cursed equilibrium is to define the refinement in the χ -virtual game introduced above. Of special interest to us is the analogue of perfect Bayesian equilibrium:

Definition 3 *σ is a χ perfectly-cursed equilibrium of G if it is a perfect Bayesian equilibrium of the χ -virtual game \overline{G}^χ .*

Perfectly-cursed equilibrium can place restrictions on beliefs off the equilibrium path, since implicit in it is the requirement that beliefs off the equilibrium path not be too extreme. In simple signaling games, for instance, when $\chi = 1$ perfect-cursedness imposes the restriction that the receiver not update her beliefs after any message, whether or not it is sent in equilibrium. There is no analog to Proposition 3 for perfectly-cursed equilibrium—the set of perfectly-cursed pooling

equilibria can depend on χ .⁹

Cursed equilibrium is the simplest way we can imagine to model players' underattentiveness to the information content of other players' actions. There are several further potential extensions and generalizations of the model that would make it potentially more realistic, but which we do not consider in this paper. We could, for example, allow for different degrees of sophistication between players, or for different degrees of sophistication for different types of a given player. While we discuss such refinements briefly in the conclusion, for the remainder of the paper we consider some key applications of our simple variant of the model, relating our results when possible to existing empirical evidence.

3 Trade

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

⁹The following sender-receiver game illustrates both the restriction that perfectly-cursed equilibrium imposes on beliefs off of the equilibrium path and the fact that not every pooling perfect Bayesian equilibrium is a perfectly-cursed equilibrium. A sender is either type t_1 or t_2 , each of which occurs with prior probability $\frac{1}{2}$; the sender knows her type, but the receiver does not. The sender chooses an action L or R . If the sender chooses R , then the game ends and both types of sender and the receiver get a payoff of 2. If the sender chooses L , then the receiver chooses between U or D . If the receiver chooses U , both types of sender and the receiver get a payoff of 4. If the receiver chooses D , then both types of sender get 0, and the receiver gets -5 if he is facing t_1 and 5 if he is facing t_2 . One perfect Bayesian equilibrium is for both types of sender to go R , and the receiver to go D if he should have the opportunity to move: going D makes sense for the receiver if he believes the deviation L comes from type t_2 of the sender with a probability of at least $\frac{9}{10}$. For sufficiently high χ , however, this is not a perfectly-cursed equilibrium. To see this, note that if $\chi > \frac{1}{5}$, the receiver's perceived payoff from facing type t_2 of sender when he plays D in \overline{G}^χ is less than 4, and therefore U dominates D . Perfectly cursed equilibrium imposes the restriction that the receiver's beliefs not be too extreme off the equilibrium path. Intuitively, this corresponds to the restriction that cursed equilibrium imposes on beliefs on the equilibrium path described in Lemma 1: if Player k thinks that type t_{-k} and t'_{-k} are both possible, then whenever $\chi > 0$ no action by players $j \neq k$ can convince Player k that she is facing t_{-k} with probability one.

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 $a : [0, 1] \rightarrow \{0, 1\}$, where $a(b) = 1$ means that he accepts the raider's offer of b . The firm's optimal strategy is clear: it sells at price b if and only if her type is less than b . Given the uniform distribution of the firm's type, therefore, the average value of the 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 the raider will get. When $\gamma < 2$, the expected net return to the raider will be negative for any positive b , so the unique Bayesian Nash equilibrium outcome involves $b = 0$. 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; this, however, is not perfectly cursed since the best response of some types of firms to a positive offer is to accept. Henceforth we limit our attention to perfectly-cursed equilibria. 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 $\bar{\sigma}_F(a)$ be the average (across types) probability that a firm plays action a . Thus, $\bar{\sigma}_F(1) = \int_0^b 1 dt + \int_b^1 0 dt = b$, 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 only accepts her offer 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.¹⁰

For $\gamma \in (2, 4)$, the raider bids too low: her payoff from bidding $b = \frac{\gamma}{4}$ is $\frac{\gamma^3 - 2\gamma^2}{32}$, which is less than $\frac{\gamma - 2}{2}$, her payoff from bidding $b = 1$. Cursedness leads to both overbidding when $\gamma < 2$ and underbidding when $\gamma > 2$ 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

¹⁰Note that even when $\gamma < 1$, the 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. Hence, despite it being common knowledge that there are no gains from trade, players trade nonetheless. While we know of no evidence on this prediction and this degree of error does not seem entirely implausible to us, it does seem somewhat unlikely.

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 , a 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

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

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

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

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 twenty 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 $[v_0, v_0 + r]$. In a χ -cursed equilibrium, the raider's optimal bid b maximizes her payoffs

$$\frac{b - v_0}{r} \left(\gamma \frac{b + v_0}{2} (1 - \chi) + \gamma \frac{2v_0 + r}{2} \chi - b \right),$$

from whence $b^* = \frac{2v_0(\gamma\chi+1)+\gamma\chi r}{4-2\gamma(1-\chi)}$.

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 1 presents both the χ -cursed equilibrium values of b and subjects' average bid \bar{b} .

Table 1: Adverse Selection (from Holt and Sherman 1994)

Curse	r	v_0	γ	$b(\chi)$	$b(\chi = 0)$	$b(\chi = 1)$	\bar{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

Holt and Sherman designed the “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 bids 3.56, while an uncursed raider bids 3. Subjects’ average bid was 3.78, slightly about the fully-cursed prediction. Finally, in the “loser’s-curse” treatment, a fully-cursed raider bids 0.81, an uncursed raider 1, and subjects 0.74. Thus, subjects’ behavior is much closer to the fully-cursed than the Bayesian-Nash prediction, although average bids depart too extremely from Bayesian Nash equilibrium to be adequately described by cursed equilibrium.

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 $\mu(\omega_1) = \mu(\omega_2) = \mu(\omega_3) = \frac{1}{3}$. Suppose that Player 2 holds an asset which pays k in state ω_k , so that the higher the state the higher the value of the asset. Each player has private information about the state of the world: Player 1 learns when the state is ω_1 , but cannot differentiate between states ω_2 and ω_3 ; Player 2 learns when the state is ω_3 , but cannot differentiate between states ω_1 and ω_2 . The information partitions $\mathcal{P}_1 = \{\{\omega_1\}, \{\omega_2, \omega_3\}\}$ and $\mathcal{P}_2 = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$ represent Player 1 and Player 2’s information, respectively; P_i is an element of Player i ’s partition \mathcal{P}_i . After each player receives her private information, Player 1 makes Player 2 an offer for the asset which Player 2 then accepts or rejects.

The only possible trade that can occur in a Bayesian Nash equilibrium of this game is the relatively meaningless one where the good is traded at price 2 in state ω_2 and neither party expects to benefit from the trade. For any $\chi \in (0, 1]$, however, trade in which a party expects to gain can occur in state ω_2 . Let $b_1 : \mathcal{P}_1 \rightarrow [1, 3]$ denote Player 1’s bidding strategy, and $a_2 : \mathcal{P}_2 \times [1, 3] \rightarrow \{0, 1\}$ denote Player 2’s acceptance strategy, where $a_2 = 1$ means Player 2 accepts Player 1’s bid. Each

player's payoff in state ω_k is k if she holds the asset after trading plus or minus any transfer she paid received or paid.

The following strategies are a cursed equilibrium with trade in state ω_2 :

$$b_1(P_1) = \begin{cases} 1 & P_1 = \{\omega_1\} \\ 2 - \frac{x}{2} & P_1 = \{\omega_2, \omega_3\}, \end{cases}$$

and

$$a_2(P_2, b_1) = \begin{cases} 1 & P_2 = \{\omega_1, \omega_2\}, b_1 \geq 2 - \frac{x}{2} \\ 0 & P_2 = \{\omega_3\} \text{ or } b_1 < 2 - \frac{x}{2}. \end{cases}$$

First note that trade cannot occur in states ω_1 or ω_3 . The most that Player 1 is willing to offer in ω_1 is 1, but because Player 2 puts positive probability on being in state ω_2 when the state is ω_2 whatever $b_1(\{\omega_2, \omega_3\})$, Player 2 rejects Player 1's offer. In ω_3 , Player 2 will accept no less than 3, but Player 1 will not offer 3 since Player 2 would accept that in state ω_2 . Now consider ω_2 . As long as $b_1(\{\omega_1\}) \neq b_1(\{\omega_2, \omega_3\})$, Player 2 thinks the probability of being in state ω_2 given he receives the bid $b_1(\{\omega_2, \omega_3\})$ is $1 - \frac{x}{2}$. Thus his expected value of the asset is $(1 - \frac{x}{2})2 + \frac{x}{2} = 2 - \frac{x}{2}$. If Player 2 accepts Player 1's offer in state ω_2 , then given that he rejects it in ω_3 , Player 1 thinks that when her offer is accepted the probability of being in ω_2 is $1 - \frac{x}{2}$, and thus the expected value of the asset is $(1 - \frac{x}{2})2 + \frac{x}{2} \cdot 3 = 2 + \frac{x}{2}$. Hence Player 1 strictly prefers to trade, and she offers $2 - \frac{x}{2}$, the lowest price at which Player 2 is willing to trade.

In this example, trade in ω_2 occurs in a cursed equilibrium because neither player sufficiently updates her beliefs about the value of the object given the willingness of the other player to trade. In the information structure given, Player 1 is overly optimistic about the value of the object based on her private information alone when it turns out that the state is ω_2 . But whereas an uncursed trader would learn from Player 2's willingness to trade at a low price that the state is ω_2 , a cursed trader remains overoptimistic that the state is ω_3 .¹¹

While trade only occurs one third of the time in the example, it is easy to see that whatever the probability of ω_2 , so long as both ω_1 and ω_3 both occur with positive probability, trade can occur. Since cursed equilibrium is consistent with speculative trade — where at least one player strictly prefers trading to not trading — with probability arbitrarily close to one, a natural question is whether it is consistent with speculative trade with probability one. It is not. To see why,

¹¹While in this trading mechanism Player 1 benefits from trade, there exist other trading mechanisms under which Player 2 gains. It is also not important to the example that both players are cursed: trade will occur in state ω_2 if only one of the two players is cursed. This follows from the fact that when Player 1 makes Player 2 an offer, Player 2 thinks that the probability of being in state ω_2 is less than one, so he will accept some offer sufficiently close to, but below, 2 when he is cursed. If Player 1 is cursed, she thinks that the probability of being in ω_2 given that her offer is accepted is less than one, and hence she is willing to offer more than 2, which Player 1 will accept.

consider again the trading mechanism described in our example where Player 1 makes Player 2 an offer, and suppose that both players are fully cursed. If Player 2 always accepts Player 1's offer, then Player 1 learns nothing about the value of the object from the fact that they are trading, and therefore she can offer no more than her expectation of the asset's value at any of her information sets. If she strictly prefers trading at some information set, then she must offer less than her expectation of the value of the object at that information set, and therefore her average offer (across all information sets) must be less than the asset's expected value. Likewise, since Player 2 is fully cursed, he infers nothing from Player 1's offer, and hence at each of his information sets he must be offered more than his expectation of the asset's value. If he strictly prefers trading at some information set, then he must be offered more than his expectation of the asset's value, and thus Player 1's average offer must exceed the expected value of the asset, a contradiction. When players are cursed, but not fully cursed, essentially the same argument applies.

4 Common-Values Auctions

In this section, we use an example to illustrate the implications of cursed equilibrium in first- and second-price auctions. Under either auction format in our example, the more cursed are bidders, the higher they bid, and when the number of bidders is sufficiently high cursed bidders suffer the winner's curse — the average winning bid exceeds the average value of the object. We show that second-price auctions raise more expected revenue than first-price auctions with cursed bidders, just as with rational bidders. However, unlike with rational bidders, as cursed bidders' information about the value of the object becomes more precise, the seller's expected revenue may fall, so a seller may have incentive to hide information about the value of the object from cursed bidders. Finally, we provide an example of a common-values auction where cursed bidders bid *less* than uncursed bidders. In the final part of this section, we discuss some of the experimental literature on common-values auctions in relation to cursed equilibrium.

In a common-values auction, the value of the object being auctioned is common but unknown to all bidders. In our example, we assume bidders receive signals that are independent and identically distributed conditional on the common value of the object. Bidders are risk neutral, and a bidder's utility from winning the auction is simply the value of the object, s , minus the price she pays, p ; her utility from losing the auction is zero. Throughout this section, we use capital letters to denote random variables and lower-case letters to denote values these random variables take on. In order

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 s is

$$(1 - \chi)s + \chi E[S|X_i = x_i] - p,$$

where x_i is the value of Bidder i 's signal about the value of the object. That is, 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.

Suppose that n bidders share a common prior on the value of the object. We follow Klemperer's (1999) example and assume that value of the object, S , is distributed uniformly on the real line, and Bidder i 's signal, X_i , is distributed uniformly on $[S - \frac{a}{2}, S + \frac{a}{2}]$ for some $a > 0$.¹² Three functions that play an important role in our analysis merit definition here: $Y_i^n(1)$ is the highest signal among bidders $j \neq i$; $r(x_i) \equiv E[S|X_i = x]$ is Bidder i 's expectation of the value of the object conditional on her signal $X_i = x$; and $v^n(x, y) \equiv E[S|X_i = x, Y_i^n(1) = y]$ is Bidder i 's expectation of the value of the object conditional on her signal being x and the highest of the other bidders' signals being y .

We say that a bidder suffers the winner's curse in a given equilibrium of a given auction if the bidder's expected surplus from entering the auction is negative; that is, the expectation of the value of the object less the price, both conditional on the event that she wins, is negative. In order that our definition apply across auction settings, we parameterize auctions by P^n , the price the winner pays when she wins the n -bidder auction; for example, in a first-price auction, P^n is the winner's bid.

Definition 4 *Bidder i suffers the winner's curse in equilibrium (b_i, b_{-i}) of the n -bidder auction if $E[(S - P^n) 1_{\{b_i(X_i) > \max_{j \neq i} b_j(X_j)\}}] < 0$.*¹³

Under our definition, a bidder suffers the winner's curse if the expected value of the object conditional on winning is less than the price conditional on winning. In a symmetric equilibrium of a symmetric model, Bidder i suffers the winner's curse if $E[S] < E[P^n]$, namely if the expected

¹²While the uniform distribution over the real line is not defined, it can be thought of as the limit of the uniform distribution on $[-K, K]$ as $K \rightarrow \infty$. When the value of the object is negative, the auction corresponds to a procurement auction where the seller pays the winner to perform some costly activity. For the purposes of the example, however, all that matters is that the bidders' beliefs about S as a function of their signals are uniform. For a more thorough analysis of Bayesian Nash equilibrium in this model, see Klemperer (1999).

¹³ $1_{\{A\}}$ is the indicator function that takes on the value one when A occurs and zero otherwise.

price exceeds the expected value of the object.¹⁴

We begin our analysis with second-price auctions, where 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 $b_i(x_i) = v^n(x_i, x_i)$ — 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 . To see that this is an equilibrium, suppose that bidders $j \neq i$ follow their proposed equilibrium strategies. A Bidder i with signal x_i who bids b_i receives a payoff of

$$\int_{x_i - a}^{b_j^{-1}(b_i)} (v^n(x_i, y) - v^n(y, y)) f_{Y_i^n(1)}(y|X_i = x_i) dy,$$

where $f_{Y_i^n(1)}(y|X_i = x_i)$ is the density of $Y_i^n(1)$ conditional on $X_i = x_i$. It is easy to show that $v^n(x_i, y)$ is increasing in x_i , which implies that the integrand is positive if and only if $x_i > y$. Hence, Bidder i 's expected utility is maximized when $b_j^{-1}(b_i) = x_i$, or $b_i = b_j(x_i)$. Intuitively, Bidder i 's bid does not affect the price she pays when she wins, only which auctions she wins. If the other bidders follow their equilibrium strategies, then the only effect of raising her bid above $v^n(x_i, x_i)$ is for Bidder i to win some auctions where $y_i^n(1) > x_i$; but in that case $v^n(x_i, y_i^n(1)) < v^n(y_i^n(1), y_i^n(1))$. In words, by raising her bid above $v^n(x_i, x_i)$, Bidder i can only win auctions she would prefer to lose. Likewise, by lowering her bid, Bidder i can only lose auctions she would prefer to win.

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

$$E\{(1 - \chi)S + \chi E[S|X_i = x_i] | X_i = x_i, Y_i^n(1) = y\} = (1 - \chi)v^n(x_i, y) + \chi r(x_i).$$

Because $r(x_i)$ and $v^n(x_i, y)$ are both increasing in x_i , the expression is increasing in x_i , and therefore we can use the same argument as Milgrom and Weber to show that $b_i(x_i) = (1 - \chi)v^n(x_i, x_i) + \chi r(x_i)$ is a χ -cursed equilibrium of the second-price auction. Here, 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. Intuitively, the second part of Bidder i 's bidding function reflects the fact that she thinks that there may be no information content in winning.

¹⁴Our definition of the winner's curse is not the only reasonable one. A more liberal definition would be that a bidder suffers the winner's curse if her expected surplus from entering the auction is less than Nash-equilibrium analysis suggests. We use our definition because it emphasizes the severity of overbidding and matches the folk wisdom that winning bids in common-values auctions tend to exceed the value of the object.

In our example, after observing signal x_i Bidder i forms posteriors that S is distributed uniformly on $[x_i - \frac{a}{2}, x_i + \frac{a}{2}]$, and so her expected value of the object conditional on her signal, $r(x_i)$, is x_i . Bidder i 's posteriors on S given that $X_i = Y_i^n(1) = x_i$ are given by

$$h^n(s|X_i = x_i, Y_i^n(1) = x_i) = \frac{\left(\frac{x_i-s}{a} + \frac{1}{2}\right)^{n-2}}{\int_{x_i-\frac{a}{2}}^{x_i+\frac{a}{2}} \left(\frac{x_i-s}{a} + \frac{1}{2}\right)^{n-2} ds} = \frac{n-1}{a} \left(\frac{x_i-s}{a} + \frac{1}{2}\right)^{n-2},$$

for $s \in [x_i - \frac{a}{2}, x_i + \frac{a}{2}]$ (and $h^n(s|X_i = x_i, Y_i^n(1) = x_i) = 0$ for $s \notin [x_i - \frac{a}{2}, x_i + \frac{a}{2}]$). Bidder i 's expectation of the value of the object conditional on her signal being both the highest and second highest on the n bidders' signals is $v^n(x_i, x_i) = x_i - \frac{a}{2} + \frac{a}{n}$. Thus, the symmetric χ -cursed equilibrium in the second-price auction is

$$b^n(x_i) = x_i - (1 - \chi)a\frac{n-2}{2n}.$$

When $n = 2$, Bayesian Nash and cursed equilibrium coincide.¹⁵ For $n \geq 3$, bids are increasing in χ for every signal value, so the seller's revenue is also increasing in χ . For $\chi < 1$, bids are decreasing in n , but the higher χ , the slower bids decrease as n increases. For a given s , the expected second-highest signal $E[Y^n(2)|S = s] = s - \frac{a}{2} + \frac{n-1}{n+1}a$, and the seller's expected revenue in the n -bidder auction is

$$E[b^n(Y^n(2))|S = s] = s - a\frac{n-1}{n(n+1)} + \chi a\frac{n-2}{2n}.$$

The seller's expected revenue is increasing in n for all χ , and as $n \rightarrow \infty$ it approaches $s + \chi\frac{a}{2} > s$, which implies that bidders suffer the winner's curse. In general, for $n > \bar{n} = \frac{\chi+2+\sqrt{9\chi^2-4\chi+4}}{2\chi}$, the seller's expected revenue exceeds s and bidders suffer the winner's curse. When $\chi = 1$, for example, $\bar{n} = 3$, meaning that bidders suffer the winner's curse whenever $n \geq 4$. As $\chi \rightarrow 0$, the χ -cursed equilibrium approaches the Bayesian Nash equilibrium where, of course, bidders never suffer the winner's curse; in this case, $\bar{n} \rightarrow \infty$.

An implication of the winner's curse is that by committing to a policy of revealing information about the value of the object, the seller may lower her expected revenue. This contrasts Bayesian-Nash analysis, where improving rational bidders' information about the value of the object mitigates bidders' fear of the winner's curse and hence intensifies the competitiveness of their bidding, raising the seller's expected revenue. In our model, as a increases, each bidder's private information about

¹⁵When $n = 2$, $r(x_i) = v^n(x_i, x_i)$, since a bidder learns nothing about the value of the object by learning that the other bidder's signal is lower than her own; intuitively, for each value of $s \in [x_i - \frac{a}{2}, x_i + \frac{a}{2}]$ is the probability that $X_j < x_i$ equal to one half. This result depends on the particular functional forms of our example, and in general Bayesian Nash and cursed equilibria can differ in two-bidder common-value auctions.

the value of the object becomes noisier, so that increasing a can be thought of as making bidders less informed. When $\chi = 0$, increasing a causes bidders to lower their bids enough that the seller's expected revenue falls. When $\chi > \bar{\chi} \equiv \frac{2(n-1)}{(n-2)(n+1)}$, however, increasing a lowers bids but *increases* the seller's revenue. For example, when $n = 4$, the seller's revenue is increasing in a whenever $\chi > \frac{3}{5}$. As $n \rightarrow \infty$, $\bar{\chi} \rightarrow 0$, so increasing a always leads to an increase in the seller's revenue, no matter slight bidders' cursedness. The winner's curse is one implication of this result. When $a = 0$, bidders know the value of the object with certainty, so the seller's expected revenue is the value of the object. For large n , increasing a increases the seller's revenue, so bidders suffer the winner's curse for $a > 0$.¹⁶

One natural question is whether the seller's expected revenue is always increasing in χ . Since $b_i(x_i) = (1 - \chi)v^n(x_i, x_i) + \chi r(x_i)$ is the χ -cursed equilibrium of the general second-price auction, the seller's expected revenue is increasing in χ whenever $E[r(Y^n(2))] > E[v^n(Y^n(2), Y^n(2))]$, namely when the expectation of the second-highest signal holder's expectation of the value of the object conditional on her signal alone is higher than the expectation of the second-highest signal holder's expectation of the value of the object conditional on her signal being the highest and second-highest. In our example, $r(x_i) = x_i$ and $v^n(x_i, x_i) = x_i - \frac{a}{2} + \frac{a}{n} < x_i$, so the seller's expected revenue does not depend on χ for $n = 2$ and is increasing in χ for $n > 2$. But consider another example where $s, x_i \in \{0, 1\}$, $\Pr[S = 0] = \Pr[S = 1] = \frac{1}{2}$, and $\Pr[X_i = 0|S = 0] = \frac{1}{2}$ and $\Pr[X_i = 0|S = 1] = 0$. When the value of the object is low, both signals are equally likely, but when the value of the object is high, the high signal occurs with probability one. In a Bayesian Nash equilibrium, a bidder with

¹⁶ As an alternative illustration that the seller may prefer withholding information from the bidders, suppose that, like the bidders, the seller receives a signal about the value of the object $Z \sim U[s - \frac{a}{2}, s + \frac{a}{2}]$. Before receiving her signal, the seller chooses between truthfully revealing and concealing her signal, whatever it is. Milgrom and Weber (1982) show that when bidders are rational the seller prefers truthful revelation. When bidders are cursed, the χ -cursed equilibrium in the auction when the seller reveals is $\tilde{b}^n(x_i, z) = (1 - \chi)\tilde{v}^n(x_i, x_i, z) + \tilde{r}(x_i, z)$. The function $\tilde{v}^n(x_i, x_i, z)$ is the analogue to $v^n(x_i, x_i)$ when the seller's signal is z , and $\tilde{r}(x_i, z)$ is the analogue to $r(x_i)$ when the seller's signal is z . It is easy to show that for all x_i and z , $\tilde{v}^n(x_i, x_i, z) = v^n(x_i, x_i)$: intuitively, if a bidder has beliefs $\mu(s)$ over $s \in [x_i - \frac{a}{2}, x_i + \frac{a}{2}]$ when her signal and the highest of other signals is x_i , then because every value of s is equally likely to generate the signal z , learning z does cause the bidder to update her beliefs. But whereas $r(x_i) = x_i$, $\tilde{r}(x_i, z) = \frac{1}{2}(x_i + z)$: a bidder's expectation of the value of the object conditional on the two signals (x_i, z) is simply their average. Thus, the seller's expected revenue in state s when she can commit to truthfully revealing her signal is

$$E[\tilde{b}^n(Y^n(2))|S = s] = s - a \frac{n-1}{n(n+1)} + \chi a \frac{n^2 + n - 4}{4n(n+1)}.$$

By concealing her signal, the seller achieves the same expected revenue as when she has no signal. When $n = 2$, the seller's expected revenue is larger than when she has no signal because $E[\tilde{r}(Y^n(2), Z)|S = s] > E[r(Y^n(2))|S = s]$, since $Y^n(2)$ is on average less than Z . When $n = 3$, the seller's expected revenue does not depend on whether she reveals her signal because $E[\tilde{r}(Y^n(2), Z)|S = s] = E[r(Y^n(2))|S = s]$, since $Y^n(2)$ is on average equal to Z . However, for $n \geq 4$, the seller's expected revenue is *lower* than when she has no signal because $E[\tilde{r}(Y^n(2), Z)|S = s] < E[r(Y^n(2))|S = s]$, since $Y^n(2)$ is on average greater than Z . Thus, with enough bidders the seller decreases her expected revenue by committing to a policy of truthfully revealing her signal.

$x_i = 0$ knows the object is worth zero, and thus $b(0) = 0$. A bidder with signal $x_i = 1$ knows that if $x_j = 0$, the object is worth zero and her payoff is zero whatever she bids. If $x_i = x_j = 1$, then the expected value of the object is $\frac{4}{5}$, and thus $b(1) = \frac{4}{5}$. When $\chi = 1$, we $b(0) = 0$ and $b(1) = \frac{2}{3}$ is a cursed equilibrium. A bidder with $x_i = 0$ knows the object is worth zero, and thus $b(0) = 0$. A bidder with $x_i = 1$ knows that the only time her bid matters is when $b_j = \frac{2}{3}$; since she thinks that b_j conveys no information about Bidder j 's signal, a bidder with $x_i = 1$ thinks the value of the object is $\frac{2}{3}$. Bidder i 's perceived expected payoff to any bid is zero, so $b(0) = 0$ and $b(1) = \frac{2}{3}$ is a fully-cursed equilibrium. Cursed and rational bidders with $x_i = 0$ both bid 0, but cursed bidders with $x_i = 1$ bid $\frac{2}{3}$, while rational bidders bid $\frac{4}{5}$. Hence, in this example, the seller's expected revenue is higher when bidders are rational than when they are cursed.¹⁷

We now turn to first price auctions, where the high bidder wins the auction and pays her bid. In a symmetric χ -cursed equilibrium of the first-price auction, a Bidder i with signal $X_i = x_i$ chooses b_i to maximize

$$\int_{\underline{x}}^{b_j^{-1}(b_i)} ((1 - \chi)v^n(x_i, y) + \chi r(x_i) - b_i) f_{Y_i^n(1)}(y|X_i = x_i),$$

where b_j is the common equilibrium bidding function of bidders $j \neq i$ and $f_{Y_i^n(1)}(y|X_i = x_i)$ is the density of $Y_i^n(1)$ conditional on $X_i = x_i$. A necessary condition for equilibrium is that

$$\frac{db^n(x_i)}{dx_i} = ((1 - \chi)v^n(x_i, x_i) + \chi r(x_i) - b^n(x_i)) \frac{f_{Y_i^n(1)}(x_i|X_i = x_i)}{F_{Y_i^n(1)}(x_i|X_i = x_i)},$$

which in our example is

$$\frac{db^n(x_i)}{dx_i} = \left(x_i - (1 - \chi)a \frac{n-2}{n} - b(x_i) \right) \frac{\int_{x_i - \frac{a}{2}}^{x_i + \frac{a}{2}} (n-1) \left(\frac{x_i-s}{a} + \frac{1}{2} \right)^{n-2} \frac{1}{a^2} ds}{\int_{x_i - \frac{a}{2}}^{x_i + \frac{a}{2}} \left(\frac{x_i-s}{a} + \frac{1}{2} \right)^{n-1} \frac{1}{a} ds},$$

which simplifies to

$$\frac{db^n(x_i)}{dx_i} = \left(x_i - (1 - \chi)a \frac{n-2}{n} - b(x_i) \right) \frac{n}{a}.$$

Hence, the symmetric χ -cursed equilibrium of the first-price auction is

$$b^n(x_i) = x_i - \frac{a}{2} + \chi a \frac{n-2}{2n}.$$

When $\chi = 0$, $b^n(x_i) = x - \frac{a}{2}$, and bids are independent of the number of bidders. When $\chi > 0$, bids increase in n . Intuitively, when $\chi = 1$, a bidder with signal x_i values the object

¹⁷We believe that "in general" revenue increases with χ . We are familiar with no experimental tests on auctions where revenues decrease with χ , which might be an additional useful test of our explanation of the winner's curse in auctions.

at $r(x_i) = x_i$, so the auction is one of private, but correlated, values. As n increases, bidders bid more because they face increased competition. For a given s , the expected highest signal $E[Y^n(1)|S = s] = s + \frac{n-1}{n+1}\frac{a}{2}$, and the seller's expected revenue in the n -bidder auction is

$$E[b^n(Y^n(1))|S = s] = s - \frac{a}{n+1} + a\chi\frac{n-2}{2n}.$$

Like in the second-price auction, the seller's expected revenue is increasing in n . Bidders suffer the winner's curse when $n \geq \bar{n} \equiv \frac{2+\chi+\sqrt{9\chi^2+4\chi+4}}{2\chi}$. When $\chi = 1$, $\bar{n} \approx 3.5$, so bidders suffer the winner's curse whenever $n \geq 4$. In a second-price auction when $\chi = 1$ bidders also suffer the winner's curse whenever $n \geq 4$. When $\chi = \frac{1}{2}$, bidders suffer the winner's curse in a first-price auction when $n \geq 6$, while they suffer the winner's curse in a second-price auction when $n \geq 5$. This difference reflects the fact that in a cursed equilibrium, as in a Bayesian Nash equilibrium, the second-price auction raises more expected revenue than the first-price auction. In a cursed equilibrium, bids are decreasing in a , just as in the rational case. When $\chi > \bar{\chi} \equiv \frac{2n}{(n-2)(n+1)}$, again the seller's expected revenue is *increasing* in a . Just as in second-price auctions, in a first-price auction with a large number of bidders $\bar{\chi}$ is close to zero, so the seller's expected revenue in large auctions is increasing in a , as long as bidders are not completely rational.

Rather than analyze more general implications of cursed equilibrium in auctions, we conclude this section by relating our analysis above to some of the large body of experimental evidence. In an early experiment, Bazerman and Samuelson (1983) auctioned off jars of coins to student subjects. In each auction, subjects could see the jar being auctioned, but did not know how many coins it contained. The highest bidder paid her bid and received the paper-dollar equivalent of the coins in the jar. Subjects also guessed how many coins each jar contained, and the subject whose guess was closest to the true value won a cash prize. Whereas all of the jars actually contained \$8.00, the average winning bid was \$10.01. However, the subjects' average estimate of the money in the jar was only \$5.13. Even though the subjects were on average too pessimistic about the value of the money in the jars, they suffered the winner's curse, presumably because those with high bids bid close to their estimates, rather than tempering their bids.

Kagel and Levin (1986) test a model nearly identical to our example above: the value of the object is distributed uniformly over $[\underline{s}, \bar{s}]$, and each bidder i receives a signal $X_i \sim U[s - \frac{a}{2}, s + \frac{a}{2}]$, when s is the value of the good. In a first-price auction, the χ -cursed equilibrium of this auction is

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

for $x_i \in [\underline{s} + \frac{a}{2}, \bar{s} - \frac{a}{2}]$ and $z_i = \exp\left(-\frac{n(x_i - (\underline{s} + \frac{a}{2}))}{a}\right)$.¹⁸ This bidding function differs from that derived in the example above only by the final term, which becomes small as x_i increases above $\underline{s} + \frac{a}{2}$; we ignore this final term in most of our discussion below. Table 2 summarizes Kagel and Levin’s data on a large series of auctions (some series aggregate auctions with different values of a , n , \underline{s} , and \bar{s}).

Table 2: Common-Values Auctions (Kagel and Levin 1986)

n	Obs	$\pi(\chi = 0)$	$\pi(\chi = 1)$	$\bar{\pi}$	χ
3 – 4	31	9.51	3.25	3.73	0.92
4	18	4.99	−0.75	4.61	0.07
4	14	6.51	−3.82	7.53	0
4	19	8.56	−0.12	5.83	0.31
4	23	6.38	−2.24	1.70	0.54
5	18	5.19	−1.90	2.89	0.32
5 – 7	11	3.65	−5.19	−2.92	0.74
6	18	4.70	−10.11	1.89	0.19
6 – 7	25	4.78	−10.03	−0.23	0.34
7	26	5.25	−8.07	−0.41	0.42
7	14	5.03	−11.04	−2.74	0.48

The first column reports the number of bidders in each auction series. The second column reports the number of auctions in each series. The third and fourth columns present the average equilibrium profits — the average winning bid less the average value of the object — when $\chi = 0$ and $\chi = 1$, respectively.¹⁹ The fourth column contains subjects’ actual average profits. The fifth column provides estimates of χ given the subjects’ behavior.²⁰

Kagel and Levin’s data are broadly consistent with positive χ , but not $\chi = 1$. In every auction series but one, subjects’ profits lie between the Bayesian-Nash and fully-cursed predictions.

¹⁸The interested reader can derive the cursed bidding function by using the fact that

$$b^n(x) = (1 - \chi)b_{\chi=0}^n(x) + \chi b_{\chi=1}^n(x),$$

where $b_{\chi=0}^n(x)$ is the Bayesian-Nash bidding function and $b_{\chi=1}^n(x)$ is the fully-cursed bidding function, both of which are presented in Kagel and Levin (1986).

¹⁹For $x_i > \bar{s} - \frac{a}{2}$, since the bidding function for rational bidders cannot be solved analytically, Kagel and Levin approximate it using the bid function for $x_i \in [\underline{s} + \frac{a}{2}, \bar{s} - \frac{a}{2}]$. As they note, this overstates bids for high signal values and hence understates the difference between fully-cursed and rational bidding. As a result, this should bias our estimate of χ downwards.

²⁰To estimate χ , we use the fact that because the bidding function in a χ -cursed equilibrium is the χ -weighted average of the fully-rational and fully-cursed bidding functions, bidders’ profits in a χ -cursed equilibrium are also the χ -weighted average of the fully-rational and fully-cursed profits. We then use Kagel and Levin’s report of the theoretical profits for rational and fully-cursed bidders, as well as subjects’ actual profits. However, because of the way that Kagel and Levin approximate the bidding function for high signals described in the last footnote, our estimate of χ is biased downwards.

Our estimates of χ are fairly consistent across auction series, with 7 of the 11 between 0.19 and 0.54. The average value of χ (weighted by the number of observations for each values) is 0.42.²¹ In this experiment, as in many others, when the number of bidders is small, average profits are positive, but when the number of bidders is large, average profits are negative. Kagel and Levin (1986) conclude that the larger the number of bidders, the further the subjects' bids from Nash equilibrium. However, if we estimate χ separately for $n \leq 4$ and $n \geq 5$, we get estimates are 0.39 and 0.46, respectively. Thus, while χ appears to be marginally higher for large n , the fact that the two estimates are so similar suggests that subjects' cursedness is not particularly sensitive to n . As we showed in our example above, whatever χ , bidders suffer the winner's curse for large enough n .

All said, cursed equilibrium seems to fit reasonably well how profits depend on the number of bidders and the noisiness of bidders' signals, a . A further indication that the subjects exhibit cursed behavior, which (unlike Table 2) includes bids from losing bidders, comes from Kagel and Levin's (1986) estimate of the linear bidding function

$$b(x_i, a, n) = 1.00x_i - 0.74\frac{a}{2} + 0.65n,$$

(0.002) (0.03) (0.15)

where standard errors are reported below the regression coefficients.²² Because the bidding function in a cursed equilibrium is linear in neither $\frac{a}{2}$ nor in n , Kagel and Levin's estimated bidding function is somewhat hard to interpret. But the coefficient on $\frac{a}{2}$ is significantly less than the value of one predicted by Bayesian Nash equilibrium, and bids are increasing in n , rather than decreasing as

²¹Kagel and Levin's (1986) Table 3 groups the data as a function of a and n ; estimating χ from this table yields the results presented in Table 3.

Table 3: Common-Values Auctions (Kagel and Levin 1986)

n	a	$\pi(\chi = 0)$	$\pi(\chi = 1)$	$\bar{\pi}$	χ
3 - 4	24	4.52	-1.24	2.60	0.33
3 - 4	36	7.20	-0.24	3.98	0.43
3 - 4	48, 60	11.22	0.60	6.75	0.42
6 - 7	24	3.46	-3.68	-1.86	0.75
6 - 7	36	3.19	-8.51	-0.95	0.35
6 - 7	48, 60	7.12	-12.31	0.60	0.34

With one exception, all of our estimates of χ lie between 0.33 and 0.43. Whether $n \in \{3, 4\}$ or $n \in \{6, 7\}$, subjects' profits are increasing in a , as Bayesian Nash predicts. Since cursed equilibrium suggests that profits are increasing in a as long as $\chi < \bar{\chi} \equiv \frac{2n}{(n-2)(n+1)}$, the data are consistent with $\chi < 0.8$ for $n \in \{3, 4\}$. For $n \in \{6, 7\}$, however, profits are increasing in a only when $\chi < 0.35$, which appears somewhat inconsistent with our estimates of χ in this case.

²²The regression includes a subject-specific and auction-specific error term for each bid. Kagel and Levin also estimate a bidding function including an intercept term and z_i , but the estimated coefficients on these variables are insignificant.

predicted by Bayesian Nash equilibrium. Both results are consistent with cursed equilibrium. Finally, we should note that in only 71% of auctions did the high-signal holder win. In a cursed equilibrium, as in a Bayesian Nash, all of the auctions should have been won by the high-signal holder, and that they were not suggests that subjects made errors in addition to those predicted by cursed equilibrium, or that different bidders were cursed to different degrees.

Many other papers find evidence of the winner’s curse. Lind and Plott (1991) show that the winner’s curse in Kagel and Levin’s (1986) experiments is not due to any strategic effects of limited liability — the fact that subjects who lost more than some initial endowment were removed from the experiment. Dyer, Kagel and Levin (1989) report experiments using students and executives from the construction industry; all but one of the executives had experience bidding in auctions. They find that both types of subjects suffer the winner’s curse, and that the curse the curse is slightly stronger among the executives, albeit not significantly. Kagel, Levin, and Harstad (1995) test the second-price auction with the same signal structure as Kagel and Levin (1986). Again, they find that subjects’ profits are less than the Nash prediction for all n , and that bidders suffer the winner’s curse when they are sufficiently numerous. Using the same procedure for estimating χ as we did for Kagel and Levin, we estimate $\chi = 0.36$, which is fairly close to our estimate of 0.42 in the first-price auction. However, Kagel and Levin’s (1986) subjects, Kagel, Levin, and Harstad’s (1995) subjects do appear to be more cursed the larger n : when $n = 4$, $\chi = 0.18$, when $n = 5$, $\chi = 0.27$, and when $n \in \{6, 7\}$, $\chi = 0.42$.

Avery and Kagel (1997) report experimental evidence on a simple two-bidder auction where each bidder receives a signal $X_i \sim U[1, 4]$, and $u_i(x_1, x_2) = x_1 + x_2$; that is, the value of the object is simply the sum of the two bidders’ signals. The argument used above to show that $b_i(x_i) = (1 - \chi)v^n(x_i, x_i) + \chi r(x_i)$ was an equilibrium of the second-price auction where applies equally well to this model, and thus $b(x_i) = \frac{5\chi}{2} + (2 - \chi)x_i$ is the symmetric χ -cursed equilibrium of this auction. Avery and Kagel estimate the linear bidding function $b(x_i) = \alpha + \beta x_i$. Cursed equilibrium predicts that $\alpha = \frac{5}{2}\chi$ and $\beta = 2 - \chi$.

Avery and Kagel divide their subjects, who are mostly undergraduate economics students, into two groups. Inexperienced subjects have played only seven (unreported) practice auctions, and their reported data cover 18 auctions. Experienced subjects are formerly inexperienced subjects who have now participated in 25 auctions; their reported data cover 24 auctions. In this auction, cursed equilibrium makes predictions about both parameters, α and β , and but without the data there is no obvious way to estimate the χ that best fits the data. Table 4, however, compares the

average values of α and β Avery and Kagel found for inexperienced and experienced subjects to different values of χ .

Table 4: Second-Price Auctions (from Avery and Kagel 1997)

	$\chi = 0$		$\chi = 0.75$		$\chi = 1$		Actual	
Subjects	α	β	α	β	α	β	$\bar{\alpha}$	$\bar{\beta}$
Inexperienced ($n = 299$)	0	2	1.875	1.25	2.5	1	2.64 (0.68)	1.13 (0.08)
Experienced ($n = 308$)	0	3	1.875	1.25	2.5	1	1.99 (0.35)	1.34 (0.05)

From the table it can be seen that $\chi = 1$ fits inexperienced subjects' behavior well, and $\chi = 0.75$ fits experienced subjects' behavior. These estimates are roughly consistent with a couple of different formal best-fit procedures. First, we minimize the distance between $\bar{\alpha}$ and $\bar{\beta}$ and $\alpha(\chi)$ and $\beta(\chi)$ by minimizing the weighted sum

$$L(\chi, z) = z (\alpha(\chi) - \bar{\alpha})^2 + (1 - z) (\beta(\chi) - \bar{\beta})^2,$$

where $z \in (0, 1)$ is the relative weight placed on explaining α versus β . We find that $\chi = \frac{1.74+11.46z}{2+10.5z}$ for inexperienced subjects, yielding $\chi \in (0.87, 1.06)$ for $z \in (0, 1)$, and $\chi = \frac{1.28+8.67z}{2+10.5z}$ for experienced subjects, yielding $\chi \in (0.64, 0.80)$ for $z \in (0, 1)$. For $z = \frac{1}{2}$, $\chi = 1.03$ for inexperienced subjects, and $\chi = 0.77$ for experienced subjects. If instead we found the value of χ that yields $\frac{\alpha(\chi)}{\beta(\chi)}$ closest to $\frac{\bar{\alpha}}{\bar{\beta}}$, we find $\chi = 0.97$ for inexperienced subjects and $\chi = 0.74$ for experienced subjects. Thus inexperienced subjects behave very much like fully-cursed bidders, and experienced subjects appear much closer to fully-cursed than uncursed.

5 Voting

A recent rational-choice literature on voting in elections and juries assumes that people vote with a sophisticated understanding that they should predicate their votes on being pivotal. Because a voter's vote only matters when she is pivotal, she should vote as if she is pivotal, even when she suspects that she is not.²³ Being pivotal can affect a voter's preferences if she believes that other voters have private information about the proper way to vote, information that is revealed from the fact that she is pivotal. Hence, a sophisticated voter asks herself what information other

²³See Razin (2000) 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.

voters would have to make her pivotal, and then how she wants to vote when she combines that information with her own private information.

In a series of papers, Feddersen and Pesendorfer (1996, 1997, 1998) explore the implications of such sophisticated reasoning by voters. Feddersen and Pesendorfer (1996) study a variant of this reasoning in which uninformed voters 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. In this section we apply cursed equilibrium to the model developed in Feddersen and Pesendorfer (1998) of a jury that must decide whether to convict a defendant of some crime. We discuss some general implications of cursedness in this model, as well as how our results fit the findings of McKelvey’s and Palfrey’s (1998) experimental test of the model.

A jury of size $M \geq 2$ must decide whether to convict some defendant of some crime. Let ω_G be the state of the world where the defendant is guilty, and 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 an innocent vote. Let $\sigma_k : \{\gamma, \iota\} \rightarrow \Delta\{g, i\}$ 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, n_G^{-i} denote the number of jurors $j \neq i$ who vote guilty, and $n_I = M - n_G$ denote the number of jurors who vote innocent. Let $a \in \{A, C\}$ be the outcome of the jury process, where A denotes acquit and C 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 $n_G \geq N$.

All jurors share the preferences

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

where $q \in (0, 1)$ is a parameter measuring the voters’ trade-offs associated with either convicting the innocent and acquitting the guilty. The higher 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 only matters if she is pivotal, it only matters if exactly $N - 1$ other jurors cast guilty votes. Thus a 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

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

where $\sigma \in [0, 1)$. Under strategy σ_k , Juror k votes guilty with probability one when she receives a guilty signal, and votes guilty with probability σ when she receives an innocent signal. Feddersen and Pesendorfer show 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 why this leads to $\sigma > 0$, note for instance that when $\frac{N}{M} > \theta$ voters realize that even when the person is guilty they typically will not convict him based on guilty votes alone; if all those with innocent signals were to vote innocent, then a person with one of those innocent signals should realize that if she is pivotal it is almost surely the case that the defendant is 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.

In order for there to be a mixed-strategy χ -cursed equilibrium of the form described above with $\sigma \in (0, 1)$, a juror must be indifferent between voting guilty and innocent when she gets a ι signal. The expected utilities for Juror i with an innocent signal, ι , from each of her two possible votes are

$$\begin{aligned} u(g|\iota) &= \Pr [n_G^{-i} \neq N - 1 | \iota] \cdot u(g | n_G^{-i} \neq N - 1) + \Pr [n_G^{-i} = N - 1 | \iota] \cdot V_g(\sigma) \\ u(i|\iota) &= \Pr [n_G^{-i} \neq N - 1 | \iota] \cdot u(i | n_G^{-i} \neq N - 1) + \Pr [n_G^{-i} = N - 1 | \iota] \cdot V_i(\sigma), \end{aligned}$$

where $V_g(\sigma)$ and $V_i(\sigma)$ are the juror's *perceived* payoffs from voting guilty or innocent if she is pivotal and receives an innocent signal. Because $u(g | n_G^{-i} \neq N - 1) = u(i | n_G^{-i} \neq N - 1)$ — a

voter only cares about her vote when it is pivotal — $u(g | \iota) = u(i | \iota)$ if and only if $V_g(\sigma) = V_i(\sigma)$.

From this,

$$\sigma^* = \max \left\{ 0, \frac{\theta z - (1 - \theta)}{\theta - (1 - \theta)z} \right\}, \text{ where } z = \left(\frac{1 - q - \theta\chi}{q - (1 - \theta)\chi} \right)^{\frac{1}{N-1}} \left(\frac{1 - \theta}{\theta} \right)^{\frac{M-N+1}{N-1}}.^{24}$$

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

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

and $\sigma^* = 0$ otherwise.

That is, when χ is small, then those with innocent signals vote guilty with positive probability when N is close enough to M , just as Feddersen and Pesendorfer found. More generally, cursed equilibrium shares many features of Bayesian Nash equilibrium. For example, when χ is sufficiently small, jurors with innocent signals sometimes vote guilty. Various comparative statics hold irrespective of χ . For all 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 the jurors are to vote guilty.

Although partially-cursed jurors may vote strategically, they underinfer one another's information when they condition their votes on being pivotal. This affects their voting strategy, and hence the extent to which voting is efficient — the likelihood that an innocent defendant is acquitted and a guilty defendant convicted. The formula above shows that $\frac{\partial \sigma^*}{\partial \chi} \leq 0$, meaning that the more cursed are jurors, the less likely are jurors with innocent signals to vote guilty. 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 of the striking results in Feddersen and Pesendorfer (1998) is that fixing the number of jurors, M , the probability of convicting an innocent person may increase as the number of guilty

²⁴To determine the equilibrium values of $V_g(\sigma)$ and $V_i(\sigma)$, we first define

$$\begin{aligned} A(\sigma) &= \Pr \left[\iota, n_G^{-i} = N - 1 \mid \omega_G \right] = \binom{M-1}{N-1} (\theta + (1-\theta)\sigma)^{N-1} ((1-\theta)(1-\sigma))^{M-N} (1-\theta) \\ B(\sigma) &= \Pr \left[\iota, n_G^{-i} = N - 1 \mid \omega_I \right] = \binom{M-1}{N-1} ((1-\theta) + \theta\sigma)^{N-1} ((\theta(1-\sigma))^{M-N} \theta, \end{aligned}$$

the probability that Juror i receives an innocent signal and is pivotal in the guilty and innocent states, respectively. A cursed juror with an innocent signal who knows she is pivotal believes that the defendant is guilty with probability $P_g \equiv (1 - \chi) \frac{A(\sigma)}{A(\sigma) + B(\sigma)} + \chi(1 - \theta)$ and innocent with probability $P_i \equiv (1 - \chi) \frac{B(\sigma)}{A(\sigma) + B(\sigma)} + \chi\theta$. By observing that $V_g(\sigma) = P_i \cdot (-q)$ and $V_i(\sigma) = P_g \cdot (q - 1)$, $V_g(\sigma) = V_i(\sigma)$, we get the result presented in the text.

votes needed for conviction, N , increases; this is because the probability with which a juror with an innocent signal votes guilty may increase 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 .

Feddersen and Pesendorfer (1998) characterize the likelihood of acquitting a guilty defendant and convicting an innocent defendant under the unanimity rule when the size of the jury becomes arbitrarily large. Under the unanimity rule, an innocent defendant is convicted with probability $\Pr[C|\omega_I] = [(1 - \theta) + \theta\sigma^*]^M$ and a guilty defendant with $\Pr[A|\omega_G] = 1 - [(1 - \theta) + \theta\sigma^*]^M$. When $\chi < \frac{1-q}{\theta}$,

$$\begin{aligned}\lim_{M=N \rightarrow \infty} \Pr[C|\omega_I] &= \left(\frac{(1-q)(1-\theta) - \theta(1-\theta)\chi}{q\theta - \theta(1-\theta)\chi} \right)^{\frac{\theta}{2\theta-1}} \\ \lim_{M=N \rightarrow \infty} \Pr[A|\omega_G] &= 1 - \left(\frac{(1-q)(1-\theta) - \theta(1-\theta)\chi}{q\theta - \theta(1-\theta)\chi} \right)^{\frac{1-\theta}{2\theta-1}}.\end{aligned}$$

$\Pr[C|\omega_I]$ is decreasing in χ , and $\Pr[A|\omega_G]$ is increasing in χ : cursedness decreases the probability of convicting an innocent defendant and increases the probability of acquitting a guilty defendant. When $\chi > \frac{1-q}{\theta}$, $\Pr[C|\omega_I] = 0$ and $\Pr[A|\omega_G] = 1$; that is, sufficiently cursed jurors vote their signals, so the defendant is never convicted.²⁵

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 one half and one. While winning an election typically requires a majority of votes, passing a proposition often requires a supermajority such as $\frac{2}{3}$.²⁶ To consider the role of cursedness in such contexts, we consider the limit as M becomes very large and when $N = kM$, where $k > \frac{1}{2}$ is a fixed parameter representing the percentage of guilty votes needed to “convict”. In this case,

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

²⁵When $\chi = 0$, $\lim_{M=N \rightarrow \infty} \Pr[C|\omega_I] = \left(\frac{(1-q)(1-\theta)}{q\theta}\right)^{\frac{\theta}{2\theta-1}}$ and $\lim_{M=N \rightarrow \infty} \Pr[A|\omega_G] = 1 - \left(\frac{(1-q)(1-\theta)}{q\theta}\right)^{\frac{1-\theta}{2\theta-1}}$, which coincides with the results in Feddersen and Pesendorfer (1998).

²⁶In a multi-candidate race with only two viable candidates, requiring a majority to avoid a run-off amounts *de facto* to requiring a super-majority.

When $\chi < \frac{1-q}{\theta}$, neither χ nor q affects the equilibrium proportion of guilty votes in the limit.²⁷ But both χ and q help to 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 $\chi < \frac{1-q}{\theta}$.²⁸ 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 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 does *not* depend 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

²⁷The intuition for this independence from χ and q depends on the fact that in a mixed-strategy equilibrium a voter must be indifferent between voting innocent and voting guilty when she is pivotal and has an innocent signal. Recall from an earlier footnote that a cursed voter with an innocent signal who knows he is pivotal believes that the defendant is guilty with probability $P_g \equiv (1 - \chi) \frac{A(\sigma)}{A(\sigma) + B(\sigma)} + \chi(1 - \theta)$ and innocent with probability $P_i \equiv (1 - \chi) \frac{B(\sigma)}{A(\sigma) + B(\sigma)} + \chi\theta$, where $A(\sigma)$ and $B(\sigma)$ are the actual probabilities that a voter receives an innocent signal and is pivotal in the two states. Since the voter is indifferent between voting guilty and innocent only if $P_i \cdot (-q) = P_g \cdot (q - 1)$, these equations imply that σ must be such that $\frac{A(\sigma)}{B(\sigma)} \in (0, 1)$. Intuitively, if the voter's perceived probabilities — and, hence, the actual probabilities — of the two states were not of the same order of magnitude, then she would strictly prefer voting innocent or to guilty. Because $\frac{A(\sigma)}{B(\sigma)} = \left[\frac{\theta + (1 - \theta)\sigma}{(1 - \theta) + \theta\sigma} \right]^{kM - 1} \left[\frac{(1 - \theta)}{\theta} \right]^{(1 - k)M + 1}$,

$$\lim_{N \equiv kM \rightarrow \infty} \frac{A(\sigma)}{B(\sigma)} = \lim_{N \equiv kM \rightarrow \infty} \left[\left(\frac{\theta + (1 - \theta)\sigma}{(1 - \theta) + \theta\sigma} \right)^k \left(\frac{(1 - \theta)}{\theta} \right)^{(1 - k)} \right]^M.$$

That is, $\frac{A(\sigma)}{B(\sigma)}$ is a likelihood function describing the relative probability that a voter with an innocent signal is pivotal in each of the two states; if the fraction being raised to the power M does not equal 1, then in the limit $\frac{A(\sigma)}{B(\sigma)}$ is either infinite or zero. In fact, σ^* is the value of σ that such that $\lim_{N \equiv kM \rightarrow \infty} \frac{A(\sigma)}{B(\sigma)} = 1$.

²⁸This can be proven by noting that when $\theta = 1$, $(1 - \theta) + \theta\sigma^* = 0$ and $\theta + (1 - \theta)\sigma^* = 1$, and that when $\theta \searrow \frac{1}{2}$, $(1 - \theta) + \theta\sigma^* = \theta + (1 - \theta)\sigma^* = k$. The result is then established by showing that $(1 - \theta) + \theta\sigma^*$ is strictly decreasing in θ and $\theta + (1 - \theta)\sigma^*$ is strictly increasing.

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 McKelvey and Palfrey (1998), who study the laboratory behavior of students at Cal Tech. Subjects were assigned randomly to groups with either 3 or 6 members. Each group was assigned with equal probability to one of two urns, the “innocent” urn with 7 innocent balls and three guilty balls, or the “guilty” urn with 3 innocent balls and 7 guilty balls.²⁹ Subjects did not know to which urn their group had been assigned, but each subject privately and independently drew a ball at random (sequentially with replacement) from her group’s urn. After observing her ball, each subject voted either innocent or guilty. McKelvey and Palfrey’s experiment corresponds to parameter values of $\mu(\omega_G) = \mu(\omega_I) = .5$, $q = .5$, and $\theta = .7$ in the model outlined above. Different groups faced different rules determining how their votes were aggregated into a decision. There were four different conditions: unanimous and majority rules in 3- and 6-person juries. That is, they ran four different combinations of M and N : $(N, M) = (2, 3)$, $(3, 3)$, $(4, 6)$, and $(6, 6)$.³⁰ Subjects received 50 cents if their group’s decision matched their urn and 5 cents if it did not.

McKelvey and Palfrey (1998) analyze their data using *quantal-response equilibrium* (QRE), both to test how well Feddersen and Pesendorfer’s model explains behavior and to test how well QRE explains subjects’ errors. Quantal-response equilibrium posits that the subjects make mistakes with some frequency, making a greater number of errors the less costly are those errors, but otherwise play a best response to other subjects’ behavior, taking into account the errors these others are making.

In principle, one could define a χ -cursed quantal-response equilibrium by combining a cursed misunderstanding of the relationship between actions and signals with the error structure embedded in quantal-response equilibrium. While we do not conduct this (complicated) analysis, we use their results to make some crude attempt to say whether cursedness adds any explanatory power the results.³¹ Subjects faced eight situations—each of the four voting rules, and each of the two possible

²⁹We follow the authors in using the language of “guilty” and “innocent” in describing the experiment, although the actual states described to the subjects were the more neutral terms “Red” and “Blue”.

³⁰For each of these four cases, McKelvey and Palfrey ran an additional condition, which we do not analyze, in which subjects conducted a non-binding “straw poll” before voting.

³¹Subjects’ exhibit a statistically-unlikely greater tendency to vote guilty on an innocent signal than innocent on a guilty signal, even in the three-voter, majority-rule case where subjects simply should vote their signals. From this, we infer that there was some “spillage” among conditions in which subjects primed to vote one way in the asymmetric

signals.³² In six of the eight contingencies—in all cases where the observed signal is γ , and in the two majority-rule cases where the signal is ι —predicted behavior does not depend on χ . The first two lines of Table 5 supply some statistics on the two cases where behavior depends on χ —the voters who have received innocent signals on three- and six-member unanimous juries .

Table 5: Jury Voting (from McKelvey and Palfrey 1998)

M	N	s	$\sigma^*(0)$	$\sigma^*(1)$	$\bar{\sigma}$	σ^{**}	Errors	Cost per Error
3	3	ι	.31	.00	.36	.00	36%	.02
6	6	ι	.65	.00	.48	1.00	52%	.03
Majority/ ι			.00	.00	.14	.00	14%	.14
All/ γ			1.00	1.00	.95	1.00	5%	.20

Columns 4 and 5 report the shares of voters in the Bayesian Nash and cursed equilibria who should vote guilty, and Column 6 shows the percentage of subjects, $\bar{\sigma}$, who actually voted guilty. As can be seen, too many people voted guilty in the three-person anonymous case—the *opposite* of the error predicted by cursedness. On the other hand, too few people voted guilty in the six-person unanimous case, consistent with cursedness. Column 7 indicates how each individual subject should have voted had she known how others were voting. Given that too many subjects voted guilty in the three-person case, the optimal strategy for an individual voter would be to vote innocent for sure; given that too many were voting innocent in the six-person case, the optimal strategy would be to vote guilty for sure. Hence, Column 7 shows that 36% of subjects were voting erroneously in the direction opposite of cursedness in the three-person case, and 52% of subjects were voting erroneously in the direction predicted by cursedness in the six-person case, suggesting that subjects were more prone to cursed errors than uncursed errors. For further comparison, the third row of Table 5 lists together the other two innocent-signal conditions, indicating that only 14% of subjects make errors in these cases. The fourth row shows average behavior by subjects getting guilty signals in the four conditions, indicating that only 5% of subjects vote incorrectly in these cases.

Subjects make more errors in the one case where those errors are “cursed” than in any other case. However, these error rates is complicated by the fact that some errors are costlier than others. Column 9 shows the expected cost of each error, measured in terms of how much each error lowers the expected likelihood of reaching the correct verdict. Since the expected cost of voting incorrectly in the conditions represented in the last two rows is much higher than in the first two conditions, cases did so in the symmetric conditions as well (or subjects were biased towards voting for red balls over blue).

³²The number of votes taking place in each of this eight situations varied between 143 and 202; in the two rows of Table 5 where we average across conditions we take the simple average of the conditions rather than weighting by the number of subjects.

the lower number of errors may merely reflect their costliness rather than the uncursed nature of the error. Yet the expected cost of voting innocent in the six-person case is greater than the expected cost of those voting guilty in the three-person case, which suggests that the greater number of these cursed errors cannot be fully explained by their low cost.³³

An alternative method of estimating subjects' cursedness in this experiment is by computing the maximum-likelihood estimate of χ under the maintained hypothesis that almost nobody makes any error *except* cursedness. In this case, the huge number of subjects voting incorrectly in the six cases where cursedness should not affect behavior are merely "flukes", and we look for the χ that best fits subjects' behavior in the two cases where the mixed-strategy played depends on χ . In this case, the maximum-likelihood value of χ is .10.³⁴ Because neither method of estimating χ is very satisfying, and neither method yields a very high estimate of χ , we conclude that whatever support McKelvey and Palfrey's voting data provide for cursedness is very weak.

6 Signaling

In this section, we briefly apply cursed equilibrium to two different signaling contexts, starting with classical simple signaling games. Because it causes the receiver to infer less from signals than she should, it is natural to suppose 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 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 \tilde{p}_l and \tilde{p}_h represent the receiver's beliefs about the probability that the sender is type g following signals l and h . After observing the signal the receiver chooses an action $a \in [0, 1]$ and has utility function

³³If instead we compared the expected cost of the error conditional on being pivotal, the difference would be more dramatic: 19% vs. 6%, rather than 3% vs. 2%.

³⁴This solves the maximum-likelihood from the observed data in these two cases:

$$\max_{\chi} L(\chi) \equiv \frac{\binom{186}{67} (\sigma^*(M = N = 3, \chi))^{67} (1 - \sigma^*(M = N = 3, \chi))^{121}}{\binom{186}{89} (\sigma^*(M = N = 6, \chi))^{89} (1 - \sigma^*(M = N = 6, \chi))^{97}}$$

$u(a, g) = -(1 - a)^2$ and $u(a, b) = -a^2$.³⁵ Hence, a receiver with beliefs \tilde{p} about the sender's type maximizes his expected utility $-\tilde{p}(1 - a)^2 - (1 - \tilde{p})a^2$ by choosing $a = \tilde{p}_l$ and $a = \tilde{p}_h$ following signals $e = l$ and $e = h$.

We assume that there is a continuous, increasing function $f : \mathbb{R} \rightarrow \mathbb{R}$ and real numbers $c_b > c_g > 0$ such that (presented in reduced form that integrates the receiver's optimal response of $a = \tilde{p}$) $u_b = u_g = f(\tilde{p}_l)$ is the payoff to both types of sender if the signal l is sent, while $u_b = f(\tilde{p}_h) - c_b$ and $u_g = f(\tilde{p}_h) - c_g$ 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 $\tilde{p}_l = \frac{\chi}{2}$ and $\tilde{p}_h = 1 - \frac{\chi}{2}$. 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 signaling can occur when the receiver is fully cursed. Intuitively, no sender would send a costly signal that would not affect the receiver's beliefs.³⁶

While fully-cursed receivers always destroy the potential for signaling, however, less extreme cursedness can *create* the potential for successful signaling. 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.

We now turn to an example of signaling that we call a “the revelation game,” modeled after politicians who feel constrained not to lie to voters, but who do not feel constrained to reveal the

³⁵The 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.

³⁶While 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 .

full truth. In the 1999-2000 American presidential campaign, candidate George W. Bush has said that he has never had an extramarital affair, and that he has not used cocaine in the past 25 years. But he refuses to say whether he used cocaine more than 25 years ago. Especially since Governor Bush volunteered the precise number 25, 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 $[0, 1]$. The sender chooses a message $m \in [0, 1] \cup \{S\}$: she either announces that she is some type in $[0, 1]$ or chooses S , meaning she remains silent. After observing the sender's message, the receiver forms beliefs about the sender's type; let $P^m(t)$ be the receiver's beliefs about the probability that the sender's type is less than t following the message m . We assume that the receiver picks an action $a(m) \in [0, 1]$ to maximize the expectation of his payoff $-(a(m) - t)^2$. This means that 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 $m \in \{t, S\}$ and $-a(m) - c$ if $m \notin \{t, S\}$. Hence, she wants the receiver's beliefs to be as low as possible, but she pays a cost of c if she misreports her type. We assume that $c > 1$, so no sender ever has incentive to misreport her type.

The most plausible Bayesian Nash equilibrium in this game is that all types reveal themselves fully.³⁷ 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 $\chi \frac{1}{2} + (1 - \chi) \left(\frac{1}{2} + \frac{r}{2} \right) = \frac{1}{2} + (1 - \chi) \frac{r}{2}$, 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

³⁷In fact, if we defined the game such that the sender *cannot* misreport her type, then this would be the unique perfect Bayesian equilibrium. But because we have not defined cursed equilibrium for games where a player's action space depends on her type, we could not apply cursed equilibrium to this game. In the conclusion, we discuss some of the problems that accompany cursed equilibrium in such games. In the game as we have defined it, there are other perfect Bayesian equilibria. One is that each type of sender chooses the action S , and the receiver chooses $a(S) = \frac{1}{2}$ and $a(m) = 1$ for $m \neq S$. In this equilibrium, no sender reveals her type because the receiver "punishes" any announcement of the sender's types with the extreme action $a = 1$. This strategy does not survive other refinements such as iterated weak dominance or the intuitive criterion—the receiver should not believe that the message m could be sent by any type of sender other than $t = m$ since any other type $t' \neq m$ could do better by announcing either $m = t'$ or $m = S$, whatever the receiver's continuation strategy.

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.

We conclude with an experiment by Forsythe, Isaac, and Palfrey (1989) that provides some evidence for cursedness in a game very similar to the revelation game. Their game is inspired by the American film industry, where movie distributors auction the rights to show films to movie-theater operators. Forsythe, Isaac, and Palfrey report that over 90% of films are auctioned off before they are shot. Theater owners dislike this practice, possibly because they suffer a winner's curse on movies auctioned before being shot. Distributors privately informed about the quality of their films pre-production who are obliged to reveal quality post-production may auction off bad movies before production and good ones after, much as in the revelation game good types reveal while bad types conceal. In Forsythe, Isaac and Palfrey's (1989) 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 on $\{1, 2, \dots, 125\}$. Each seller knew the value of her object, but the bidders did not. The sellers chose whether to reveal the value of their objects to the bidders or conceal them; a seller who revealed her value had to do so truthfully. Following this, the objects were auctioned to the bidders using first-price auctions, where each buyer bid on each of the sellers' items. 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, while sellers with objects valued less than $r = \frac{125\chi+1}{1+\chi}$ conceal their values. Intuitively, low-value sellers conceal the value of their objects because cursed bidders mistakenly think that some high-value sellers conceal, causing them to bid too high for objects whose values are concealed. When $\chi = 0$, all sellers (except possibly ones 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 6 summarizes the data.

Table 6: Revelation Game (from Forsythe, Isaac, and Palfrey 1989)

Group	Sellers	Conceal	Value Conceal	Bid Conceal	χ
All	240	85	31	39	0.44
Experienced	120	32	23	27	0.27
Experienced*	72	12	11	19	0.17

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. For sellers who concealed, the average value of their objects was 31, but the average winning bid was 39. Hence buyers suffered a winner’s curse, paying an average of 8 cents more than the value of the objects they won. There are two natural ways to estimate χ from the data, one from the average winning bid and one from the average value of the objects concealed. Consider the data from the first row and suppose that the sellers follow a cutoff strategy—revealing when their values were high and concealing otherwise. Then because the average value of the sellers’ objects is 31, sellers would be revealing when their objects were worth more than 61 and concealing otherwise. Then since in equilibrium $61 = r = \frac{125+\chi}{1+\chi}$, $\chi \simeq 0.83$; that is, sellers who thought that bidders were cursed with $\chi = 0.83$ would reveal with values over 61 and conceal otherwise. If sellers were following this strategy, however, then cursed bidders would bid an average of 61, far more than the 39 that they actually bid. The other method of estimating χ is to assume that cursed bidders believe that sellers follow a cutoff strategy, and estimate χ from the average winning bid. In this case, $\chi \simeq 0.44$. In Table 6, we present estimates of χ using the second of these methods since it corresponds better to bidder behavior as well as to seller behavior with the notable exception of a few outlying high-value objects whose values were concealed in early rounds of the experiment: 68 of 88 sellers (77%) with values less than 39 concealed, and 17 of 152 (11%) of sellers with values more than 39 concealed; by contrast, only 74 of the 128 (58%) of the sellers with values less than 61 and 11 of the 112 (10%) sellers with values greater than 61 concealed. Thus, given bidders’ behavior, 37 of 240 sellers (15%) made mistakes: the 20 with objects worth less than 39 who revealed, and the 17 with objects worth more than 39 who concealed. In other words, those sellers whose objects were worth less than 39 and revealed had objects with significantly higher valuations than those sellers with objects worth less than 39 who concealed.³⁸

³⁸In fact, seller behavior is better described by a simple step rule: sellers with the lowest valued concealed; sellers with intermediate values concealed half of the time; and sellers with high values revealed. 43 of 44 sellers with objects valued less than 25 concealed, and the average value of their objects was 14; the value of the object of the lone seller who revealed was 16. 31 of 58 sellers (53%) with values between 25 and 49 concealed, and the average value of their

The second row of the table describes only those subjects who have already participated in 10 trials or more — trials 11 to 16 for the one group that participated in 16 trials, and trials 11 to 22 for the other two groups. 27 percent of sellers concealed the value of their objects; the average value of these sellers' objects was 23, and the average winning bid was 27. Thus even experienced bidders suffered a winner's curse, albeit half of what it was in the aggregate data. Again, if cursed bidders in these final rounds believed that sellers followed a cutoff strategy, concealing if their value was less than r and revealing otherwise, then from the data $r = 27$, so $\chi \simeq 0.27$. Seller behavior is very close to this as 20 of the 22 sellers (91%) with objects worth 27 or less concealed, and 12 of 98 sellers (12%) of sellers with objects worth more than 27 concealed.

Finally, because in one of the three groups a single subject won 16 of the 20 of the auctions where the seller concealed her value, bidding an average of 35, the third row excludes this subject's group from the pool of experienced subjects. This time, only 17% of sellers concealed the value of their objects. The winner's curse is larger than for experienced subjects, as the average value of sellers' objects was 11, while the average winning bid was 19. Thus the winner's curse was larger. For this group, $\chi \simeq 0.17$.

7 Discussion and Conclusion

We believe that cursed equilibrium can provide insight in many additional domains. One is in organizational and sequential decision-making, where we believe that cursedness may capture a form of exaggerated fear that some parties may have of putting other parties in charge of decisions, underappreciating the fact that unanticipated future decisions by others may be based on unanticipated information. Consider, for instance, a grand jury that must decide whether to indict some defendant of a crime. If the defendant is indicted, the case proceeds to trial where a jury hears the evidence and decides whether to convict the defendant. In this case, a sufficiently cursed grand jury that is not yet convinced of the defendant's guilt may be too reluctant to indict. This is because it fears that the jury will convict when the defendant is innocent, even though it should realize that the jury only convicts if it has strong evidence that the defendant is guilty. Similarly, principals in organizations may be reluctant to delegate decisions even to parties whose interests coincide with

objects was 34; the average value of the 27 who revealed was 36. Finally, 11 of the 138 sellers (8%) with objects worth at least 50 concealed (8%), and the average value of their objects was 88; the average value of the 127 who revealed was 86. The fact that within each of these groups the average value of the objects of sellers who concealed roughly equals the average value of the objects of those sellers who revealed suggests that sellers' decisions to reveal did not depend on their objects' values.

their own out of fear that the these other parties would make different decisions than they would, underappreciating how often those different decisions reflect superior information.³⁹

Many applications of cursed equilibrium point to limitations and problems with the solution concept as we have defined it, and we conclude by discussing some of these shortcomings and possible extensions of the solution concept. One limitation of cursed equilibrium is that we have only defined it in games where each player’s action space is independent of her type. In games without such independence, cursed equilibrium should presumably be defined such that players do not assign positive probability to a type playing an action that is impossible for that type. A problem with this approach, however, is that a cursed equilibrium in the game where an action is impossible 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 a modified definition in which we assumed 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 take that action. Developing a new concept incorporating this notion seems important both intuitively and for practical application, but would raise new problems such as determining how to measure and compare how irrational a given action is for different types of a player, and precisely how to restrict beliefs as a function of the degree of irrationality associated with a rule. And the enterprise 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.

Perhaps a more urgent direction for developing the idea of cursed equilibrium concerns a more important limitation to our current definition. The notion of cursed equilibrium is meant to capture a general intuition that people tend to underappreciate the relationship between others’ actions and the information these others have at the time they take those actions. Yet our formal definition

³⁹To illustrate with a simple principal-agent model, let $\Omega = \{\omega_1, \omega_2\}$ be the set of possible states of the world, where the principal and agent share the common prior $\mu(\omega_1) = \mu(\omega_2) = \frac{1}{2}$. If the principal invests in an experiment, the agent learns which is the true state, otherwise the agent learns nothing about the state. Once the agent has received his information, he chooses an action $a \in \{a_1, a_2, a_3\}$. The principal and the agent share the common payoffs $u(a_1, \omega_1) = u(a_2, \omega_2) = 2$, $u(a_1, \omega_2) = u(a_2, \omega_1) = -2$, and $u(a_3, \omega) = 1$ for each ω . That is, the agent attempts to match action a_i to state ω_i , while the safe action a_3 pays one in each state. Then if the agent learns the true state, he matches his action to the state, earning a payoff of 2. If the agent does not learn the state, he chooses a_1 , earning a payoff of 1. A rational principal therefore prefers that the agent learn the true state. But if the principal is cursed, her perceived payoff from performing the experiment is $(1 - \chi)2 + \chi(\frac{1}{2}2 + \frac{1}{2}(-2)) = 2(1 - \chi)$, which exceeds one only when $\chi < \frac{1}{2}$. Thus a sufficiently cursed principal prefers that the agent not learn the true state and hence take the safe action, because she inappropriately fears that the “risky” action following the experiment might mismatch the state.

makes an artificial distinction between private information represented by a type space in a Bayesian game and private information that is not represented by the type space. 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 1 might take that Player 2 observes but Player 3 does not. We hope to move towards a more complete notion of cursed equilibrium which allows for "cursedness" over more general types of unobservable information that others have.⁴⁰

Many other generalizations of cursed equilibrium seem important to add more realistic variation in the degree of "cursedness" in different situations. For instance, we find it intuitive that players are less likely to ignore the informational content of given actions by other players when they have not actually observed those actions than when they have; observing actions seems likely to induce more strategic interpretations. This might imply that the reactions by players to the observed actions in certain sequential games are "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 is not affected by whether voters vote sequentially or simultaneously. While we believe the same equivalence holds with cursed equilibrium as we have defined in this paper, a better model may have a cursed voter understand the relationship between other voters' signals and votes better when she can observe their votes than when she cannot, leading to more rational voting in the sequential than in the simultaneous-move voting procedure.

A final generalization of cursed equilibrium, manifestly needed to more tightly fit the data, is to allow different players to be cursed to different extents. Such heterogeneity is the natural interpretation of many of the experiments cited above, for instance; while we believe that in many cases behavior was usefully characterized by positing a uniform $\chi > 0$ across subjects, the behavior would be even better described by allowing for heterogeneity.⁴¹

⁴⁰The different treatment of "exogenous" and "endogenous" private information seems not only intuitively and psychologically wrong to us, but 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 an alternative Bayesian game where a fictitious player is added who takes actions observable by the privately-informed player, our definition of cursed equilibrium is not robust.

⁴¹In fact, we suspect that in some circumstances heterogenous cursedness may lead to some qualitatively-different predictions than homogenous cursedness.

8 Appendix

Proof of Lemma 1 From Bayes' Rule,

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

Proof of Proposition 1 Consider the alternative game (A, T, p, \bar{u}^χ) , where (A, T, p) are all the same, but u is replaced by

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

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. σ is a Bayesian Nash equilibrium of \bar{G}^χ if for each Player k and each type $t_k \in T_k$, and each a_k^* such that $\sigma_k(a_k^*|t_k) > 0$,

$$\begin{aligned}a_k^* \in \arg \max_{a_k \in A_k} & \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k} \in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) \bar{u}_k^\chi(a_k, a_{-k}; t_k, t_{-k}) \\ &= (1-\chi) \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k} \in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) u_k(a_k, a_{-k}; t_k, t_{-k}) \\ & \quad + \chi \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k} \in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}).\end{aligned}$$

But

$$\begin{aligned}& \chi \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k} \in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}) \\ &= \chi \sum_{a_{-k} \in A_{-k}} \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \sigma_{-k}(a_{-k}|t_{-k}) \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}) \\ &= \chi \sum_{a_{-k} \in A_{-k}} \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}) \bar{\sigma}_{-k}(a_{-k}|t_k) \\ &= \chi \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \sum_{a_{-k} \in A_{-k}} \bar{\sigma}_{-k}(a_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}),\end{aligned}$$

and hence

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

Thus if σ is a Bayesian Nash equilibrium of \overline{G}^χ , it is also a cursed equilibrium of G . Because \overline{G}^χ is finite, it has a Bayesian Nash equilibrium, and so G has a cursed equilibrium. \square

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}^χ 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}^χ , which by Proposition 1 is the set of χ -cursed equilibria of G . Hence it suffices to show that

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

The second expression can be rewritten

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

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

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

as desired. \square

Proof of Proposition 3 Suppose that σ is 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

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

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

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

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 $\chi \in [0, 1]$. \square

References

- [1] Akerlof, George. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84, August 1970: 488-500.
- [2] Athey, Susan. "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information," mimeo February 1997, MIT.
- [3] Avery, Christopher and John H. Kagel. "Second-Price Auctions with Asymmetric Payoffs: An Experimental Investigation." *Journal of Economics & Management Strategy* 6(3), Fall 1997: 573-603.
- [4] Ball, Sheryl.B.; Bazerman, Max H.; and John S. Carroll. "An Evaluation of Learning in the Bilateral Winner's Curse," *Organizational Behavior and Human Decision Processes* 48:1-22.
- [5] Bazerman, Max H. and William F. Samuelson. "I Won the Auction But I Don't Want the Prize," *Journal of Conflict Resolution* 27(4), December 1983: 618-634.
- [6] Dekel, Eddie and Michele Piccione. "Sequential Voting Procedures in Symmetric Binary Elections," *Journal of Political Economy* 108(1), February, 2000: 34-55.
- [7] Dyer, Douglas Kagel, John H. and Dan Levin. "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis," *Economic Journal* 99, March 1989: 108-115.
- [8] Feddersen, Timothy and Wolfgang Pesendorfer. "The Swing Voter's Curse," *American Economic Review* 86(3), June, 1996: 408-424.
- [9] Feddersen, Timothy and Wolfgang Pesendorfer. "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica* 65(5), 1997: 1029-1058.
- [10] Feddersen, Timothy and Wolfgang Pesendorfer. "Convicting the Innocent: the Inferiority of Unanimous Jury Verdicts under Strategic Voting," *American Political Science Review*, 90(1), 1998: 23-35.
- [11] Forsythe, Robert; Isaac, R. Mark; and Thomas R. Palfrey. "Theories and Tests of 'Blind Bidding' in Sealed-Bid Auctions," *RAND Journal of Economics* 20(2), Summer 1989: 214-238.

- [12] Garvin, Susan and John H. Kagel. "Learning in Common Values Auctions: Some Initial Observations," *Journal of Economic Behavior & Organization* 25, 1994: 351-372.
- [13] Jacobsen, Ben; Potters, Jan; Schram, Arthur; van Winden, Frans; and Jörgen Wit. "(In)accuracy of a European Political Stock Market: The Influence of Common Value Structures," *European Economic Review* 44, 2000: 205-230.
- [14] Klemperer, Paul. "Auction Theory: A Guide to the Literature," *Journal of Economic Surveys* 13(3), 1999: 227-260.
- [15] Holt, Charles A. and Roger Sherman, "The Loser's Curse." *American Economic Review* 84(3), June 1994: 642-652.
- [16] Kagel, John H. "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, John H. Kagel and Alvin E. Roth, eds, Princeton NJ: Princeton University Press: 1995.
- [17] Kagel, John H. and Dan Levin. "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76(5), December 1986: 894-920.
- [18] Levin, Dan; Kagel, John H.; and Jean-Francois Richard. "Revenue Effects and Processing in English Common Value Auctions," *American Economic Review* 86(3), June 1996: 442-460.
- [19] Lind, Barry and Charles R. Plott. "The Winner's Curse: Experiments with Buyers and Sellers," *American Economic Review* 81(1), 1991: 335-346.
- [20] McKelvey, Richard D. and Thomas R. Palfrey. "An Experimental Study of Jury Decisions." Mimeo, California Institute of Technology, July 2, 1998.
- [21] Milgrom, Paul and John Roberts. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica* 58(6), November 1990: 1255-1277.
- [22] Milgrom, Paul and Nancy Stokey. "Information, Trade, and Common Knowledge," *Journal of Economic Theory* 26: 17-27.
- [23] Milgrom, Paul R. and Robert J. Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50(5), September 1982: 1089-1122.
- [24] Potters, Jan and Jörgen Wit. "Bets and Bids: Favorite-Longshot Bias and Winner's Curse," CentER Discussion Paper 9604, Tilburg University, June 1995.

- [25] Razin, Ronny. "Swing Voters and Mandate Makers: Voting when Margin of Victory Matters," Mimeo, Princeton University, April 1999.
- [26] Samuelson, William F. and Bazerman, Max H. "The Winner's Curse in Bilateral Negotiations." In *Research in Experimental Economics, vol. 3*, Vernon L. Smith, ed., Greenwich, CT: JAI Press, 1985.
- [27] Thaler, Richard H. "The Winner's Curse," *Journal of Economic Perspectives* 2(1), 1988:191-202.