Notes on Inferential Naivety in Games

WORK IN PROGRESS FOR BACKGROUND MATERIAL FOR SEMINARS IN BRESCIA AND BOCCONI; NOT FOR GENERAL DISTRIBUTION

Erik Eyster
Department of Economics
LSE

and

Matthew Rabin
Department of Economics
University of California—Berkeley

May 5, 2007

Abstract

We develop a solution concept where a player in a strategic setting may fail to fully appreciate the informational sophistication of other players. In signalling models, our framework incorporates the idea that players are *guileless*—play as if others do not make inferences from their own behavior—or are *credulous*—play as if others do not try to manipulate their own beliefs. In information-cascade models, it captures the idea that players may naively believe that another player's actions solely reflect that player's private information. This can lead players to become very confident about the true state of the world even in environments where rational players would never become certain, as well as to herd on incorrect actions, even in environments where fully rational players never incorrectly herd.

Keywords: Init?, Didn't I?
JEL Classification: B49

Acknowledgments: We thank Zack Grossman for valuable research assistance, Marco Ottaviani, Peter Sorensen, and Georg Weizsäcker for helpful conversations, and seminar participants at Berkeley, Birkbeck/UCL, Davis, ESSET Gerzensee, Harvard/MIT, Lausanne, and Wash U for helpful comments, and the NSF and [[[XXX]]] for financial support. Eyster thanks Berkeley for its hospitality while part of this research was undertaken.

Contact: e.eyster@lse.ac.uk; rabin@econ.berkeley.edu

1 Introduction

Although game theory emphasizes players' sophisticated extraction and manipulation of information in strategic contexts, people may handle information in more limited ways. For instance, a player might be *inferentially naive* by failing to fully appreciate how another player makes inferences from her own actions or those of some other player, how that other player attempts to manipulate her own beliefs, or how that other player uses her private information more generally. Receivers in signaling contexts may underappreciate how senders attempt to manipulate their beliefs for strategic advantage; senders in signaling games may underappreciate how receivers make inferences from their actions. People in social-herding situations may not realize that the actions of somebody they observe may reflect not only that person's private information but also inferences which that person has made from the actions of still other people. In this paper, we develop a formal model that captures these and other forms of inferential naivety within a unifying framework that can be applied to any game with private information. The extreme form of our solution concept roughly corresponds to the idea that each player knows that all players use their private information but fails to recognize that other players also know this. We apply the solution concept to simple signaling games and, more extensively, to social learning (information cascade) models.

In Section 2 we define our solution concept for standard finite Bayesian games where players' private information is represented by their "types," prove existence, and connect our solution concept to Bayesian Nash equilibrium. Our model builds off of Eyster and Rabin's (2005) notion of *cursed equilibrium*, which assumes that players under-infer the informational content of one another's play. Formally, in a cursed equilibrium each player incorrectly believes that with positive probability each type profile of the other players plays the average mixed action profile that all types of the other players play, rather than that type profile's true action profile. Players choose their actions to maximize their expected utilities given their types and these incorrect beliefs about other players' equilibrium strategies. The parameter $\chi \in [0, 1]$ represents the probability that a player assigns to other players' playing their average mixed action profile rather than their true, type-contingent strategy and characterizes the extent to which a player is "cursed". The parameter value $\chi = 0$ corresponds to the fully rational Bayesian Nash equilibrium, whereas $\chi = 1$ corresponds to the case where each player assumes no connection whatsoever between other players' actions and their types, a "fully cursed equilibrium".

This paper uses cursed equilibrium as a point of departure to model players who are naive in their informational inferences. Players who are fully cursed make no informational inferences from one another's play. A player who believes that other players do not make informational inferences may therefore believe that these other players play their parts of a fully-cursed equilibrium. We define *inferentially naive information transmission* (INIT) as play in which uncursed players may exaggerate the likelihood that other players are cursed. In this extreme version of our solution concept, each player mistakenly believes that all other players play some fully-cursed equilibrium and best responds to those (mistaken) beliefs. Since players in a fully-cursed equilibrium incorrectly believe that there is no relationship between other players' actions and information, INIT play captures the notion that each player believes she uniquely understands the information content in other players' actions. In N-player games, we define ν-INIT play for $\nu \equiv (\nu_1, \ldots, \nu_N) \in [0,1]^N$ to be play in which each Player k believes that with probability ν_k all other players play a fully cursed equilibrium and with probability $1 - \nu_k$ that these other players play their true, type-contingent strategies. We extend the concept to allow for the possibility that players *are* to some degree cursed by defining χ-cursed, ν-INIT play, where each Player k also believes that with probability χ_k others' strategies do not depend upon their information.[1]

This construction founded upon cursed equilibrium provides a mechanism for capturing a notion of inferential naivety across a broad array of situations. Because fully cursed players fail to appreciate that other players' actions reflect those other players' information, it predicts not only that players infer nothing from others' behavior but also that they do not engage in costly signaling. By having players mistakenly believe that other players are fully cursed, INIT play captures inferential naivety because players believe that others do not use information the way they in fact do.

In Section 3, we apply INIT to simple signaling games, where it corresponds in intuitive ways to

[1]This method of combining partial naivety with partial cursedness is admittedly somewhat *ad hoc*, and our applications emphasize the pure $\chi = 0, \nu = 1$ case. Nevertheless, we hope that our more general definitions should prove useful in empirical implementation of the concepts developed in this paper.

different errors players may make. A sender engaging in INIT play is "guileless" in the sense that she believes that the receiver infers nothing from her play; as a result, her choice of action reflects pure preferences over actions rather than any attempt to manipulate the receiver's beliefs. When the receiver alone is INIT, then in ν-INIT play he is "credulous" in the sense that he believes that the sender does not attempt to manipulate his beliefs; he interprets the sender's actions at face value rather than as attempts by the sender to manipulate his beliefs.

We also consider nearly-cheap-talk games, where the sender uses a natural language that includes her type space and has some slight preference for announcing her true type. We show how credulity leads receivers to believe the sender's message, which accords with evidence by Malmendier and Shanthikumar (2006) that small investors ignore their financial advisors' incentives in touting one form of asset over another.

In Section 4 we apply INIT to our main example, information-cascade models, where we develop two types of results. First, in the standard cascade model with two actions, INIT play gives rise to beliefs that are more extreme than in Bayesian Nash equilibrium: players who herd on A eventually come to believe that A is the true state of the world. In INIT play, each player believes that predecessors' actions depend upon their signals alone, and so each A choice is interpreted as another piece of evidence in favour of A. This model fits experimental evidence about beliefs directly as well as evidence that longer cascades get broken less often than shorter ones. Second, we introduce a variant of the standard cascade model with two possible states of the world and a continuum of both signals and actions; unlike the standard model, the action space is rich enough to reveal players' beliefs. In the model, rational players choose actions that strictly increase in and therefore reveal their posteriors; with probability one, players' actions and beliefs converge to the true state of the world. Not so in INIT play. While actions and beliefs in INIT play do converge, they do not always converge to the true state. Intuitively, INIT players collectively overweight early signals because each Player k believes that Player $k-1$'s action depends upon $k-1$'s signal alone, while in fact it depends upon the actions (and signals) of all previous players. When the first player takes an action that signals a belief that A is likely to be the true state, and the second player takes the very same action, then the third player interprets this as two bits of evidence in favour of A rather than the single bit that it truly is. In this way, INIT play gives rise to a kind of social confirmation bias, where early hypotheses about the state of the world are too unlikely to be rejected. Based on the idea that actions and beliefs converge faster to the true than to the false state, we show that when beliefs are slow to converge players most likely believe in the wrong state. If for many periods, T, the first T players all judge the likelihood of A to be between 0.6 and 0.9—they remain far from certain that A is the true state—then in fact B is more likely to be the true state.

We conclude the paper in Section 5 with a discussion of possible extensions of the notion of cursed equilibrium that might cope with these problems, as well as some possible further economic applications of the principles developed in the paper.

2 A Model of Informational Naivety

In this section, we formally define our various solution concepts, prove their existence for a broad class of finite Bayesian games, and develop some general principles and results.

Consider a finite Bayesian Game, $G = (A_1, \ldots, A_N; T_0, T_1, \ldots, T_N; p; u_1, \ldots, u_N)$, played by players $k \in \{1, ..., N\}$. 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. $T \equiv T_0 \times T_1 \times ... \times T_N$ is the set of type profiles, and p is the prior probability distribution over T, which we assume is common to all players. Player k's payoff function $u_k : A \times T \to \mathbb{R}$ depends on all players' actions $A \equiv A_1 \times ... \times A_N$ and their types. A (mixed) strategy σ_k for Player k specifies a probability distribution over actions for each type: $\sigma_k : T_k \to \triangle A_k$. Let $\sigma_k(a_k|t_k)$ be the probability that type t_k plays action a_k, and let $u \equiv (u_1, ..., u_N)$.

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 $T_{-k} \equiv \underset{j \neq k}{\times} T_j$ of other players (including nature) given her own type $t_k \in T_k$. Let $A_{-k} \equiv \underset{j \neq 0,k}{\times} A_j$

be the set of action profiles for players $j \neq k$ (excluding nature, who takes no action), and $\sigma_{-k} : T_{-k} \rightarrow \underset{j \neq 0,k}{\times} \triangle A_j$ be a strategy of Player k's opponents, where $\sigma_{-k}(a_{-k}|t_{-k})$ is the probability that type $t_{-k} \in T_{-k}$ plays action profile a_{-k} under strategy $\sigma_{-k}(t_{-k})$.

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

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

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

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

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

When Player k is of type t_k, $\overline{\sigma}_{-k}(a_{-k}|t_k)$ is the probability that players $j \neq k$ play action profile a_{-k} when they follow strategy σ_{-k}. A player who (mistakenly) believes that each type profile of the other players plays the same mixed action profile believes that the other players are playing $\overline{\sigma}_{-k}(\cdot|t_k)$ whenever they play $\sigma_{-k}(a_{-k}|t_{-k})$. Note that $\overline{\sigma}_{-k}(a_{-k}|t_k)$ depends on t_k, so different types of Player k have different beliefs about the average action of players $j \neq k$. Let $\overline{\sigma}_{-k}(t_k) : T_{-k} \rightarrow \underset{j \neq 0,k}{\times} \triangle A_j$ denote t_k's beliefs about the average strategy of players $j \neq k$, namely $\overline{\sigma}_{-k}(t_k)$ is the strategy players $j \neq k$ would play if each type profile t_{-k} played a_{-k} with probability $\overline{\sigma}_{-k}(a_{-k}|t_k)$.

From this, we define a heterogenously cursed equilibrium, defined with respect to a profile of parameters $\chi_k \in [0, 1]$, where χ_k measures the degree to which Player k misperceives the correlation between the other players' actions and types. Let $f(\chi_1, \chi_2, ..., \chi_N; t_1, t_2, ..., t_N)$ be a probability mass function that gives the probability of any cursedness and type profile for the players.

Player k's strategy, $\sigma_k(\cdot|t_k, \chi_k)$, gives a probability distribution over her actions a_k as a function of her type, t_k, and cursedness, χ_k. Given players' cursed-contingent strategies, for each *combination* of type and cursedness of each player, we are interested in her perception of the average strategy of the other players, averaged over their types. For each $t_k \in T_k$ and $\chi_k \in [0,1]$ such that $f(t_k, t_{-k}; \chi_k, \chi_{-k}) > 0$ for some χ_{-k} and t_{-k}, let $g(t_{-k}, \chi_{-k}|t_k, \chi_k) \equiv \frac{f(t_k,t_{-k};\chi_k,\chi_{-k})}{\sum_{t'_{-k},\chi'_{-k}} f(t_k,t'_{-k};\chi_k,\chi'_{-k})}$, the probability of (t_{-k}, χ_{-k}) conditional on Player k's type and cursedness, (t_k, χ_k). We define

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

When Player k is of type t_k and cursedness χ_k, $\overline{\sigma}_{-k}(a_{-k}|t_k, \chi_k)$ is the probability that players $j \neq k$ play action profile a_{-k} when they follow strategy σ_{-k}. Its interpretation coincides with $\overline{\sigma}_{-k}(a_{-k}|t_k)$ in the text: a player who (mistakenly) believes that each type profile of the other players plays the same mixed action profile believes that the others play $\overline{\sigma}_{-k}(\cdot|t_k, \chi_k)$ whenever they in fact play $\sigma_{-k}(a_{-k}|t_{-k})$. Note $\overline{\sigma}_{-k}(\cdot|t_k, \chi_k)$ does not depend upon χ_k through the degree of misprediction implicit in $\overline{\sigma}_{-k}(\cdot|t_k, \chi_k)$; $\overline{\sigma}_{-k}(\cdot|t_k, \chi_k)$ represents the beliefs of a Player k who believes that others players' strategies do not depend on their types. Rather, χ_k affects $\overline{\sigma}_{-k}(\cdot|t_k, \chi_k)$ only insofar as it may be correlated with χ_{-k}, which may affect the strategies of Players $-k$.

Definition 2 *A mixed-strategy profile σ is a f-*generalized cursed equilibrium *if for each k, $t_k \in T_k$ and each χ_k in the support of f, and each a_k^* such that $\sigma_k(a_k^*|t_k, \chi_k) > 0$,*

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

A special case that plays an important role in this paper is where each Player k has $\chi_k = 1$ with probability one, which we call a *fully cursed equilibrium*. In a fully cursed equilibrium, each player correctly predicts the distribution over other players' actions but incorrectly believes that these other players' actions do not depend upon their types. Hence, fully cursed players mistakenly believe that other players' actions convey no information about these other players' types.

When all players share the same value of χ, χ-cursed equilibrium corresponds to a Bayesian Nash equilibrium in a modified game where the players' payoffs for each action and type profile are the χ-weighted average of their actual payoffs and their average payoffs for that action profile averaged over other players' types. As a consequence, such homogenous cursed equilibria exist in all finite Bayesian games described in this section; in particular, fully-cursed equilibrium exists.[2] Let $\Sigma^{FCE}(G)$ denote the set of fully-cursed equilibria of the game G. For $\sigma^{FCE} \in \Sigma^{FCE}(G)$, define $\overline{\sigma}_{-k}^{FCE}(a_{-k}|t_k)$ from σ^{FCE} as above.

In a fully cursed equilibrium, all players act as if the actions of others contain no information. Suppose, however, that each player feels that she uniquely understands the informational content of other players' actions—and that all other players are fully cursed.

Definition 3 *A mixed-strategy profile $\hat{\sigma}$ is* inferentially-naive information transmission (INIT) *if for each k, there exists $\sigma^{FCE} \in \Sigma^{FCE}$ such that for each $t_k \in T_k$ and each a_k^* for which $\hat{\sigma}_k(a_k^*|t_k) >$ 0,*

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

In INIT play, each player best responds to the beliefs that all other players play their parts in some fully-cursed equilibrium. In games with a unique fully-cursed equilibrium, all players share the same beliefs about which fully-cursed equilibrium the other players are playing. More generally, however, INIT play does not impose any restriction on the joint distribution of players' beliefs beyond that each believes all the others to be playing in the manner prescribed by some fully-cursed equilibrium.

Definition 4 *For $\nu = (\nu_1, \ldots, \nu_N) \in [0,1]^N$, the mixed-strategy profile $\hat{\sigma}$ is ν-*INIT *play if for each Player k, there exists $\sigma^{FCE} \in \Sigma^{FCE}$ such that for each $t_k \in T_k$, and each a_k^* such that $\hat{\sigma}_k(a_k^*|t_k) > 0$,*

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

In ν-INIT play, each player k best responds to beliefs that all other players play the ν_k-weighted combination of their actual and fully-cursed strategies. This captures the idea that players attribute to one another some, but less than perfect, inferential sophistication. Unlike in a generalized-cursed equilibrium, INIT players may have incorrect beliefs about the probability distribution over other

[2] Eyster and Rabin (2005) prove existence proof of homogenous cursed equilibrim along these lines. Existence of heterogenous cursed equilibrium follows from a fixed-point argument much like the one used here to establish existence of INIT play.

players' actions. Note that our formulation imposes the restriction that Player k is as naive about Player i's strategy as about Player j's.

Rather than best respond to the ν_k-weighted average of other players' actual play and their play in the game's fully-cursed equilibrium, players may cursedly best respond to these beliefs. Define $\overline{\hat{\sigma}}_{-k}(a_{-k}|t_k)$ to be the cursed perception of $\hat{\sigma}_{-k}$, constructed as above.

Definition 5 *For $(\chi;\nu) = (\chi_1,\ldots,\chi_N;\nu_1,\ldots,\nu_N) \in [0,1]^{2N}$, the mixed-strategy profile $\hat{\sigma}$ is* χ-cursed, ν-INIT *play if for each Player k, there exists $\sigma^{FCE} \in \Sigma^{FCE}$ such that for each $t_k \in T_k$ and each a_k^* such that $\hat{\sigma}_k(a_k^*|t_k) > 0, a_k^* \in \arg\max_{a_k \in A_k}$*

$$\sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)\cdot \sum_{a_{-k}\in A_{-k}} \big[\nu_k \big((1-\chi_k)\sigma^{FCE}_{-k}(a_{-k}|t_{-k}) + \chi_k \overline{\sigma}^{FCE}_{-k}(a_{-k}|t_k)\big) + (1-\nu_k)\big((1-\chi_k)\hat{\sigma}_{-k}(a_{-k}|t_{-k}) + \chi_k$$

In cursed INIT play, players both mispredict the distribution of other players' actions—thinking they play closer to a fully-cursed equilibrium than they do actually—as well as underestimates the degree to which those other players' actions depend upon their types.

Proposition 1 *If $G = (A, T, p, u)$ is a finite Bayesian game, then for each $(\chi,\nu) \in [0,1]^{2N}$, G has a χ-cursed, ν-*INIT *strategy profile.*

Proof For any strategy profile σ, define $BR_{\chi,\nu}(\sigma)$ to be the set of χ-cursed best-responses to the ν-weighted combination of σ and some fully-cursed equilibrium σ^{FCE}, whose existence is proved by Eyster and Rabin (2005). It is easy to verify that $BR_{\chi,\nu}(\sigma)$ has closed graph and is non-empty, compact-, and convex-valued for each σ; hence, Kakutani's Fixed-Point Theorem implies the existence of a fixed point, which is a χ-cursed, ν-INIT strategy profile. **Q.E.D.**

Two refinements of INIT prove useful for applications. Loosely speaking, both impose restrictions on the beliefs that players hold over $\Sigma^{FCE}(G)$.

Definition 6 *For $(\chi,\nu) = (\chi_1,\ldots,\chi_N;\nu_1,\ldots,\nu_N) \in [0,1]^{2N}$, the mixed-strategy profile $\hat{\sigma}$ is* homogenous-belief *χ-cursed, ν-*INIT *play if there exists $\sigma^{FCE} \in \Sigma^{FCE}$ such that for each k, each $t_k \in T_k$, and each a_k^* for which $\hat{\sigma}_k(a_k^*|t_k) > 0, a_k^* \in \arg\max_{a_k \in A_k}$*

$$\sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k)\cdot \sum_{a_{-k}\in A_{-k}} \big[\nu_k \big((1-\chi_k)\sigma^{FCE}_{-k}(a_{-k}|t_{-k}) + \chi_k \overline{\sigma}^{FCE}_{-k}(a_{-k}|t_k)\big) + (1-\nu_k)\big((1-\chi_k)\hat{\sigma}_{-k}(a_{-k}|t_{-k}) + \chi_k$$

In homogenous-belief INIT play, players i and j have common beliefs about the strategy of any player $k \notin \{i, j\}$. In two-player games, all INIT play satisfies this restriction.

Next, we may wish to restrict players' off-equilibrium-path actions to be best responses to some beliefs about other players' types, in the spirit of Perfect Bayesian Equilibrium. Eyster and Rabin (2005) define the perfect fully-cursed equilibria of game G to be the Perfect Bayesian equilibria of the associated game G' with the new payoffs

$$\widehat{u}_k(a; t_k, t_{-k}) \equiv \sum_{t'_{-k} \in T_{-k}} p(t'_{-k}|t_k) u_k(a; t_k, t'_{-k}).$$

Each type t_k of Player k's payoff in G' from the action profile a when the other players' types are t_{-k} is simply the weighted average of Player k's payoff from a in G, averaged over other players' types $t'_{-k} \in T_{-k}$ from t_k's perspective. Letting $\Sigma^{PFCE}(G)$ be the set of perfect fully-cursed equilibria in G, INIT play is then said to be perfect if we can replace $\Sigma^{FCE}(G)$ with $\Sigma^{PFCE}(G)$ in Definition 3. In a perfect fully-cursed equilibrium, at every information set, each player plays a best response to her opponents' strategies given her priors about their types; that is, in a dynamic game, no player ever updates her beliefs about any other player's type, even off the equilibrium path. INIT players who believe that their opponents are playing a perfectly fully cursed equilibrium do update their beliefs in light of their opponents' actions. In a dynamic game, however, an INIT player may observe an opponent playing some action never played in any fully cursed equilibrium. What beliefs would an INIT player then form? When Player k observes Player j choose an action inconsistent with a fully-cursed equilibrium, we assume that Player k believes that Player j is best responding to χ_j^k-cursed beliefs, where χ_j^k is as high as possible and still consistent with Player j's action being a best response to χ_j^k-cursed beliefs for some t_j. (Since j is in fact playing a best response—$\chi_j = 0$—for some t_j, $\chi_j^k \in [0, 1]$ must exist.)

Before turning to novel predictions of INIT play, it helps to establish when INIT play coincides with Bayesian Nash equilibrium.

Proposition 2 *If the strategy profile σ is both a Bayesian Nash and fully cursed equilibrium, then for each $(\chi, \nu) = (\chi_1, \ldots, \chi_N; \nu_1, \ldots, \nu_N) \in [0, 1]^{2N}$, σ is χ-cursed, ν-INIT play.*

Since σ is both a best-response and a fully-cursed best response to itself, σ is a χ-cursed best response to σ (itself the ν-weighted combination of σ and the fully cursed equilibrium σ) and thus is INIT play. Knowing this allows us to characterize an important class of games where Bayesian Nash equilibrium and homogenous INIT play coincide.

Let $T_{-0k} \equiv \underset{i \neq 0,k}{\times} T_i$ be the set of possible types of all players $i \neq k$ excluding nature, Player 0. Let $E\left[U_k(a; t_k, t_{-k})|t_k\right]$ be Player k's expectation of her payoff when she plays action a_k and the other players play action profile a_{-k}, conditional on her type t_k; U_k is random because it may depend on t_0 or t_{-0k}. Let $E\left[U_k(a; t_k, t_{-k})|t_k, t_{-0k}\right]$ be Player k's expectation of her payoff when she plays action a_k and the other players play action profile a_{-k}, conditional on her type t_k and the other players' (excluding nature's) type t_{-0k}.

Corollary 3 *Suppose that for each Player k, each type $t_k \in T_k$, each type profile $t_{-0k} \in T_{-0k}$, and each action profile $a \in A$, $E\left[U_k(a; t_k, t_{-k})|t_k, t_{-0k}\right] = E\left[U_k(a; t_k, t_{-k})|t_k\right]$. If the strategy profile σ is a Bayesian Nash equilibrium, then it is homogenous χ-cursed, ν-INIT play for every $(\chi, \nu) \in [0,1]^{2N}$.*

The corollary follows directly from Proposition 2 coupled with Proposition 2 in Eyster and Rabin (2005), which shows that when $E\left[U_k(a; t_k, t_{-k})|t_k, t_{-0k}\right] = E\left[U_k(a; t_k, t_{-k})|t_k\right]$, fully cursed and Bayesian Nash equilibria coincide. This condition not only requires that no player's payoff be affected by another player's type, but also that no player can learn anything about her expected payoff by learning another player's type. This means essentially that (given a player's type) other players' types are uncorrelated with nature's type. This distinction is crucial in many 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 expected payoff from a profile of bids would change. Hence, the corollary 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.[3]

[3] The corollary also applies to reputation models in repeated games. When a sequence of short-run players faces a

Intuitively, cursed equilibrium differs from Bayesian Nash equilibrium only in that players have incorrect beliefs about the relationship between their opponents' actions and those opponents' types; if no player's payoff depends upon any other player's type, then such mistaken beliefs do not matter, and cursed equilibrium and Bayesian Nash equilibrium coincide. Consequently, homogenous INIT play—where all players best respond to beliefs that other players are playing some weighted combination of their actual strategies and the same fully cursed equilibrium—coincides with Bayesian Nash equilibrium.

Proposition 4 *Let G have a unique fully cursed equilibrium σ. If σ is also INIT play, then it is a Bayesian Nash equilibrium.*

Proof. Since σ is INIT play, each player plays a best response to the fully-cursed equilibrium σ, and hence σ is a Bayesian Nash equilibrium.

But note that a strategy profile that is both a Bayesian Nash equilibrium and INIT play need not be a fully cursed equilibrium. One such example is Akerlof's lemon model analysed in Eyster and Rabin (2005), where no trade is both a Bayesian Nash equilibrium and INIT play but not a fully cursed equilibrium.

3 Simple Signaling Games

In this section, we consider the implication of INIT play in simple signaling settings, which provides more intuition about INIT play in general as well as generates predictions of interest in those settings.

long-run player whose preferences are unknown, no short-run player cares directly about the long-run player's type—whether he is commited to playing a certain way—but instead cares only about which action the long-run player plays against her. We regard inability to explain a short-run player's overinferring that she is facing the commitment type from observing the commitment action as a shortcoming of our approach.

3.1 General

We start by defining simple sender-receiver signaling games. The sender has a type t_s that belongs to the finite set T_s and whose distribution is given by probability mass function $p(t_s)$; the sender knows her type but the receiver does not. The sender chooses a message m from some feasible set of messages M. A strategy for the sender σ_s specifies a mixture over messages for each type. The receiver observes m and then chooses an action a from some finite set A. A strategy for the receiver specifies a mixture over actions after each possible message. The players receive payoffs of $u_s(m, a; t_s)$ and $u_r(m, a; t_s)$; each player may care about both players' actions as well as the sender's type.

We will in our examples be concerned with predictions as a function of degree of cursedness and (especially) INITness of each player, characterizing behavior as a function of the parameters $(\chi_S, \chi_R; \nu_S, \nu_R)$. While in a particular context below we consider heterogeneity in the INITness of a particular player, we will focus for tractable illustrative purposes on cases where either the Sender or the Receiver, or both, are fully inferentially naive. In addition to the Bayesian Nash equilibrium—which is $(\chi_S, \chi_R; \nu_S, \nu_R) = (0, 0; 0, 0)$—cursed equilibrium—which is $(\chi_S, \chi_R; \nu_S, \nu_R) = (1, 1; 0, 0)$—and the pure INIT play—which is $(\chi_S, \chi_R; \nu_S, \nu_R) = (0, 0, 1, 1)$—it is worth describing two focal types of heterogenous INIT play in signaling settings:

Definition 7 *1) Play consistent with* $(\chi_S, \chi_R; \nu_S, \nu_R) = (0, 0; 1, 0)$ *is a* guileless *outcome.*
2) Play consistent with $(\chi_S, \chi_R; \nu_S, \nu_R) = (0, 0; 0, 1)$ *is a* credulous *outcome.*

For each of these, we also consider the more dynamically consistent model, and more specific model, of Perfect Bayesian Equilibrium, perfectly cursed equilibrium, perfectly INIT play, perfectly guileless play, and perfectly credulous play, as the definitions corresponding to those in the previous section.

Guileless play corresponds to the Sender acting as if the Receiver infers nothing from her message, when in fact the Receiver *does* rationally infer the Sender's expected type. In guileless play, the Sender does not try to affect the Receiver's beliefs, which the Receiver recognizes. In credulous play, the Receiver acts as if the Sender does not try to manipulate the Receiver's beliefs, whereas in fact the Sender anticipates the Receiver's inferences (given the Receiver's naive beliefs about the Sender).

To characterize equilibria, some kinds of strategies play especially important roles. For any message, we can define the set of its (generically, unique) "non-inferential best responses" as the set of probability distributions over actions that maximize the expected utility of a Receiver who maintains her prior beliefs about the Sender's type: $\overline{\Sigma}_r(m) \equiv \triangle\{a \in A | a \in \arg\max_{a' \in A} \sum p(t_s) u_r(m, a'; t_s)\}$, and define $\overline{\Sigma}_r$ as the set of message-contingent strategies by the Receiver such that $\overline{\sigma}_r \in \overline{\Sigma}_r$ iff $\overline{\sigma}_r(m) \in \overline{\Sigma}_r(m)$ for each m. Define the two players' best responses to each other's strategies by $\Sigma_s^*(\sigma_r)$ and $\Sigma_r^*(\sigma_s)$, in the usual way.[4]

Cursed equilibria, and especially perfectly cursed equilibria, are easily characterized using this notation:

Lemma 5 *1) For all $\sigma_r \in \overline{\Sigma}_r$ for all $\sigma_s \in \Sigma_s^*(\overline{\sigma}_r)$, (σ_s, σ_r) is a fully cursed equilibrium.*
2) (σ_s, σ_r) is a perfectly cursed equilibrium iff $\sigma_r \in \overline{\Sigma}_r$ and $\sigma_s \in \Sigma_s^*(\overline{\sigma}_r)$.

Lemma 5 says that any strategy in which the Receiver best responds to his priors after all messages and the Sender best responds to the Receiver's strategy is a fully cursed equilibrium—and that only such strategy pairs are perfectly cursed equilibria. Notice that, since the Receiver has no private information, the Sender in a cursed equilibrium is playing not only her perceived but also the actual best response to the Receiver's strategy. Generically, the perfectly cursed equilibrium will be unique, rendering the various forms of perfectly INIT play similarly easy to characterize. Lemma 6 helps to characterize some forms of guileless and credulous play.

Lemma 6 *1) For each $\overline{\sigma}_r \in \overline{\Sigma}_r$, $\sigma_s \in \Sigma_s^*(\overline{\sigma}_r)$, and $\sigma_r \in \Sigma_r^*(\sigma_s)$, (σ_s, σ_r) is guileless and INIT play.*
2) (σ_s, σ_r) is perfectly guileless and perfectly INIT play iff there exists $\overline{\sigma}_r \in \overline{\Sigma}_r^*$ such that $\sigma_s \in \Sigma_s^*(\overline{\sigma}_r)$ and $\sigma_r \in \Sigma_r^*(\sigma_s)$

In a guileless equilibrium, the sender best responds to beliefs that the receiver will not infer anything from any message she might send. Instead, the receiver best responds to the sender's

[4] Hence, BNE in these games are any (σ_s, σ_r) for which $\sigma_s \in \Sigma_s^*(\sigma_r)$ and $\sigma_r \in \Sigma_r^*(\sigma_s)$.

message with beliefs correctly updated from the sender's strategy. The sender is "guileless" in failing to appreciate how changing her message changes the receiver's beliefs about her type and therefore his action. The reason that guileless play is also INIT play follows from the fact that the Receiver has no private information. Hence, if the Sender is inferentially naive, she plays her fully-cursed-equilibrium strategy; best responding to the actual Sender strategy is the same as best-responding to the "mistaken" belief the Sender is playing her perfectly cursed strategy.

Lemma 7 characterizes some of the credulous equilibria.

Lemma 7 *1) For all $\overline{\sigma}_r \in \overline{\Sigma}_r$, $\sigma'_s \in \Sigma^*_s(\overline{\sigma}_r)$, for all $\sigma_r \in \Sigma^*_r(\sigma'_s)$, and $\sigma_s \in \Sigma^*_s(\sigma_r)$, (σ_s, σ_r) is credulous play.*
2) (σ_s, σ_r) is perfectly credulous play iff there exists $\overline{\sigma}_r \in \overline{\Sigma}_r$ and $\sigma'_s \in \Sigma^*_s(\overline{\sigma}_r)$ such that $\sigma_r \in \Sigma^*_r(\sigma'_s)$ and $\sigma_s \in \Sigma^*_S(\sigma_r)$.

In a credulous equilibrium, the receiver best responds to beliefs that the sender is guileless, but the sender best responds to the receiver's actual strategy. The receiver is credulous because he takes the sender's messages at face value, not recognizing the sender's guile: the receiver fails to appreciate that the sender chooses her messages to influence the receiver's beliefs.

3.2 "Pooling job-market signaling"

We begin with a simple two-type, two-action variant of the much-studied educational-signaling model.

Our emphasis in this and the two subsequent examples will be on cases where the Bayesian Nash (and hence Perfect Bayesian) equilibria will be unique.

3.3 Full-Support Job-Market Signaling

We now turn to a variant of job-market signaling that provides a crisp way of illustrating some interesting implications of cursedness, INITness, guilelessness, and credulity.

A worker of type $t \in (-\infty, +\infty)$ chooses an education level $e \in (-\infty, +\infty)$. The worker knows his type, but a potential employer only observes his education level e; the firm's priors $f(t)$ over the worker's type have full support, and we assume $E[t]$ is finite. After observing the education level, the employer chooses a wage level $w \in (-\infty, +\infty)$.

The worker has utility function $U^s(e, w, t) \equiv w - k \cdot (e - t)^2$, where $k > 0$. This implies that each type of worker has a preference, *ceteris paribus*, for choosing the education level corresponding to his type, and loses utility through better *or worse* education. The assumption that workers have non-zero ideal education levels distinguishes this model from the usual signaling model where the marginal cost of education is always positive.

The employer has preferences $U^R(e, w, t) \equiv -(w - t)^2$. This implies that an employer who infers that she is facing type t would choose wage $w = t$; it provides a reduced form of a market that pays workers according to their ability.

What are the Perfect Bayesian equilibria in this game? It can readily be confirmed that, independent of $f(t)$, there is a fully separating Bayesian Nash equilibrium:

Result 1 *For all $f(t)$ and k, a fully separating BNE always exists, with $e^{BNE}(t) = t + \frac{1}{2k}$ for all t and $w^{BNE}(e) = e - \frac{1}{2k}$ for all e. This equilibrium is also a PBE, and is the unique separating BNE.*[5]

Equilibrium involves full sorting, with an induced wage-as-function-of-type of $w(t) = t$, which (depending on the environment) can be interpreted as socially efficient hiring in the market. As usual, however, the worker chooses an inefficiently excessive education level, namely $e(t) > t$.

Although the game also has semi-separating equilibria, we focus on full separation when comparing the predictions of cursed and INIT play.[6] The corresponding predictions in this model for our three non-BNE solution concepts are as follows:

[5] Kartik, Ottaviani, and Squintani (KOS, forthcoming) analyse a related extension of Crawford and Sobel's (1982) model of cheap-talk communication. In KOS, fixing the receiver's belief about the state to be equal to the true state, the higher the state, the higher the sender's preferred message; further, this message is unique for every state. In this setting, KOS show the existence of a unique increasing and differentiable Bayesian Nash equilibrium.

[6] There are an infinity of semi-separating equilibria, separating for low types and pooling for higher types. Not all low types can pool because the intervals of t's pooling on any education level must increase with type.

Result 2 *For all priors $f(t)$, and each e and t, $e^{FCE}(t) = t$ and $w^{FCE}(e) = E[t]$ is a fully cursed and the unique perfect fully cursed equilibrium. In any fully cursed equilibrium, $w^{FCE}(e) = E[t]$ for each e in the support of $e^{FCE}(\cdot)$, which is unbounded above and below.*

In the fully cursed equilibrium, because the employer infers nothing from the worker's education level, she sets the wage equal to her unconditional expectation of the worker's quality. Knowing that education will not influence his wage, the worker chooses an education according to his intrinsic preferences and not its signaling value. The fully cursed equilibrium is more efficient that the separating PBE because the worker does not attempt to signal his type and hence chooses the efficient educational level. However, the fact that wages no longer match ability as in the separating PBE has the distributional consequence that lower-ability workers are paid more and higher-ability workers are paid less than in the separating PBE.

Result 3 *For all priors $f(t)$, for each e and t, $e^{guileless}(t) = t$ and $w^{guileless}(e) = e$ is guileless and INIT play, and is the unique perfectly guileless and unique perfectly INIT play.*

Guileless play delivers the most efficient outcome: all types of worker receive the right amount of education and get paid according to actual ability. Intuitively, guileless senders make no effort to disguise or distort their type, and (given they all have different preferences over "messages" sent) are fully revealed. If workers are inferentially naive, this outcome occurs irrespective of whether the firms are inferentially naive for reasons discussed above.

Result 4 *For all priors $f(t)$, for each e and t, $e^{credulous}(t) = t + \frac{1}{2k}$ and $w^{credulous}(e) = e$ is credulous the unique perfectly credulous play.*

Credulous play in this model nicely illustrates one point we emphasized above: INIT play, as manifested in credulous interpretation of others' behavior, is not a theory of "over-inference" *per se* in the sense of players inferring "too much" from a signal. It is a theory of mis-inference or biased inference of a particular type—a kind of "understrategic" inference. Here, in both PBE and in credulous play, the Sender's type is fully revealed by her message, and in both cases the Receiver does indeed make a fully confident inference about the Sender's type. In PBE this strong inference is the correct one, while in credulous play it is incorrect, exaggerating the quality of the worker from his education level.

These different predictions, of course, are all extreme. Of some interest is the case where there is heterogeneity among both workers and employers in degree of inferential sophistication. This turns out to be harder to analyze for general distributions $f(t)$, since there will intrinsically be a distribution of types choosing a particular education level, and a distribution of reactions by employers to that educational choice. These distributions will depend on $f(t)$.

We can, however, usefully solve for one particular case, the limiting case where $f(t)$ becomes "uniform" over the entire real line.

Suppose that there is a distribution of different values of $v_s \in [0,1]$ and $\nu_r \in [0,1]$, and assume that the degree of inferential naivety is independent between the two players, and independent of the sender's type. Let $\overline{v_s}$ be the average INITness of the Sender, and $\overline{v_r}$ the average INITness of the Receiver. It can be shown that it is generalized INIT play for each (t, v_s) combination to choose $e(t) = t + (1 - v_s)\frac{1}{2k}$ and for each v_r type of Receiver to choose $w(e) = e - (1 - v_r)(1 - \overline{v_s})\frac{1}{2k}$. Intuitively, because all types of employers have a wage response with $dw/de = 1$, the attractiveness to each worker of increasing her education does not depend on the employer's INITness but only the extent to which the worker realizes that the employer will react. The employer then responds according to his inferential naivety, and to perceived extent of guilelessness by workers. Under the uniform distribution of types and independence assumptions, the average guilelessness of workers will be independent of the observed education level.

Notice that the implied wage as a function of worker's type is the realized degree of inferential naivety is given by $w(t) = t + [(1 - v_s) - (1 - v_r)(1 - \overline{v_s})]\frac{1}{2k}$, yielding a wage averaged across degrees of inferential naivete of $\overline{w}(t) = t + \overline{v_r}(1 - \overline{v_s})\frac{1}{2k}$. This shows that workers are "overpaid" in proportion to how naive employers are and how sophisticated workers are. In cases where a worker happens to be more guileless than average and the employer happens to be pretty sophisticated, a worker of a particular type will be paid less than she would in the case where everybody is sophisticated. But it can also be seen that the general bias of inferential naivete is to overpay workers relative to the fully rational case.

4 Cascades

We now consider the role of inferential naivety in information cascades. In recent years, beginning with Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), a theoretical literature has explored the role of inference in herding behavior. In the simplest model, a group of people choose one of two options in turn. They have common preferences over the two choices but do not know which they prefer. Rather, each person receives private information about which option is better; all private information is of the same quality. In addition, each person can observe all of her predecessors' choices. In this setting, rational agents herd: if the first two choose A, then so too will the third, regardless of her private information. (Because the first has private information favouring A, and the second does also—given any natural tie-breaking rule that precludes an agent with private information favouring B from choosing A for sure —the third disregards her private information and chooses A.) And all subsequent players will follow suit. Although everyone in a rational herd may choose A, they never become certain that A is the better alternative because everyone understands that a herd forms after the first two people choose A.

In a fully cursed equilibrium, players do not appreciate the informational content of their predecessors' actions. Consequently, their actions depend only upon their own signals and not upon previous play. In INIT play, each player believes that her predecessors' actions depend upon their signals alone, and so each A choice is interpreted as another piece of evidence favouring A. Unlike rational players, each successive INIT player following a herd on A becomes more convinced that A is in fact the true state of the world.

Recently, such models have been tested extensively in the laboratory. While some predictions of the rational model have found support, several systematic discrepancies have been uncovered. Kübler and Weizsäcker (2004) show in a variant of the model (described in more detail below) that subjects' beliefs become too extreme to be well explained by rational play. Kübler and Weizsäcker (2005) report the related finding that longer cascades are more stable, intuitively because after a long string of A choices, people come to believe A more and more likely, reducing the likelihood that anyone will break the cascade by choosing B.

One widely recognized limitation of this literature is that the action space is too coarse to represent many contexts where people's actions reveal the intensity of their beliefs. Finer action spaces reduce the scope for rational herding, and in the limiting case where no two distinct beliefs give rise to the same best response, equilibrium actions converge to the best response to the true state with probability one. Nevertheless, we suspect that people may herd in such settings. In this section, we show how INIT play can lead to incorrect herding in even this informationally rich setting.

There are two possible states of the world, $\omega \in \{0,1\}$, each equally likely *ex ante*. Each player k receives the signal $s_k \in [0,1]$; signals are independent and identically distributed conditional on the state. When $\omega = 0$, signals have the density function f_0; when $\omega = 1$, they have density f_1. Each player observes her signal and the actions of all previous players before choosing an action in $[0,1]$. For simplicity, we assume that the model is symmetric: for each $s \in [0,1]$, $f_0(s) = f_1(1-s)$. In addition, we assume that f_0 is uniformly bounded and $f_0(1) = 0$, as well as that the likelihood ratio $L(s) \equiv \frac{f_0(s)}{f_1(s)}$ is continuously differentiable with image $\mathbb{R}_+$ and derivative $L'(s) < 0$. Let $a_t(a_1, \ldots, a_{t-1}; s_t)$ be the action taken by the tth player as a function of previous players' actions and her own private information, and let $a \equiv (a_1, a_2, \ldots) \in [0,1]^{\mathbb{N}}$ be the profile of all players' actions. Every Player k has payoff function $g_k(a;\omega) = -(a_k - \omega)^2$, which is maximized by setting $a_k = E[\omega|F_k]$, where F_k is all the information available to Player k.

To develop some intuition for INIT play, it helps to consider an example. Let $f_0(s) = 2(1-s)$ and $f_1(s) = 2s$, in which case the likelihood ratio $L(s) = \frac{f_0(s)}{f_1(s)} = \frac{1-s}{s}$. In INIT play, each agent thinks that all previous agents are fully cursed, whereas in reality no agent is cursed. It is easy to work out that $a_1(s_1) = s_1$: the first agent follows her signal. The second agent observes the first agent's action and his own signal before choosing an action. Because the first player would follow her signal whatever her cursedness, the second player correctly infers the first player's signal from her action and chooses

$$a_2(a_1; s_2) = \frac{a_1 s_2}{a_1 s_2 + (1-a_1)(1-s_2)} = \frac{s_1 s_2}{s_1 s_2 + (1-s_1)(1-s_2)} = \frac{1}{1 + L(s_1)L(s_2)}.$$

The third player believes the second to be fully-cursed. Because a fully-cursed Player 2 would choose $a_2(a_1; s_2) = s_2$, Player 3 mistakenly infers that Player 2's signal is $\widehat{s}_2 = a_2(a_1; s_2)$. Hence he chooses

$$\begin{aligned} a_3(a_1, a_2; s_3) &= \frac{s_1 \widehat{s}_2 s_3}{s_1 \widehat{s}_2 s_3 + (1-s_1)(1-\widehat{s}_2)(1-s_3)} \\ &= \frac{s_1^2 s_2 s_3}{s_1^2 s_2 s_3 + (1-s_1)^2(1-s_2)(1-s_3)} \\ &= \frac{1}{1+L^2(s_1)L(s_2)L(s_3)}. \end{aligned}$$

The third player's action differs from the optimal choice $a_3^*(a_1, a_2; s_3) = \frac{1}{1+L(s_1)L(s_2)L(s_3)}$ by overweighting the first signal. Intuitively, because Player 3 ignores how Player 2's action depends upon Player 1's action and hence upon Player 1's signal, Player 3 weights Player 1's signal more than she should. In general,

$$a_t = \frac{1}{1 + L(s_t) \prod_{k=1}^{t-1} L^{2^{t-1-k}}(s_k)}.$$

Early signals are overweighted so that the first signal gets half the weight (in some sense) of all signals, the second half of what remains, etc.

Because INIT play weights early signals so heavily, it seems possible that even an aribitrarily large number of players may fail to learn the true state if the first few players have inaccurate signals. The following proposition shows this to be true not just in the example but in the more general model outlined above, where the key assumption is the lack of any "magic signal" that reveals the true state with certainty.

Proposition 5: Let $f_0(t) = 2(1-t)$ and $f_1(t) = 2t$. For each $k \in \left(\frac{1}{2}, 1\right)$, there exists $m > 0$ such that $\Pr[a_t > k \; \forall t | \omega = 0] > m$.

The claim formalizes the notion that false cascades may arise in INIT play: we prove it by showing that if when $\omega = 0$ the first couple of agents receive high enough signals to take actions above k, then with positive probability no agent will ever take an action below k.

Proof To show that people might converge to the wrong actions for sure in this setting, we establish a lower bound on the value of $\Pr\left[a_1, a_2, ...a_t, ... \geq k | \omega = 0\right]$ — that is, the likelihood that all actions in the infinite series will be greater than k, given that the true state is 0. This lower bound is positive, showing that, unlike in the Bayesian Nash equilibrium, there is a chance of a cascade occurring on the wrong action.

Because $\Pr[a_1 \geq k|\omega = 0] = (1-k)^2$,

$$\Pr[a_1, a_2, ...a_t, ... \geq k|\omega = 0] > (1-k)^2 \prod_{t=2}^{\infty} \Pr[a_t \geq k|a_1 = a_2 = ... = a_{t-1} = k, \omega = 0].$$

That is, the probability that the action in period t is greater than or equal to k given that all previous actions were *at least* k is at least as high as if all the previous signals were exactly k.

The formula for the beliefs of a fully naive person t who gets signal s_t and observes $a_1 = a_2 = ... = a_{t-1} = k$ is simply $\frac{k^{t-1}s_t}{k^{t-1}s_t+(1-k)^{t-1}(1-s_t)}$. Hence, $\Pr[a_t \geq k|a_1 = a_2 = ... = a_{t-1} = k, \omega = 0]$ equals the probability that s_t satisfies $\frac{k^{t-1}s_t}{k^{t-1}s_t+(1-k)^{t-1}(1-s_t)} \geq k$ given that the state is $\omega = 0$. This can be reduced to $\Pr\left[s_t \geq \frac{1}{1+\left(\frac{k}{1-k}\right)^{t-2}} \middle| \omega = 0\right]$. It can be shown that $\Pr\left[s_t \geq \frac{1}{1+\left(\frac{k}{1-k}\right)^{t-2}} \middle| \omega = 0\right] = \left(\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}}\right)^2$, which using our earlier formula implies that $\Pr[a_1, a_2, ...a_t, ... \geq k|\omega = 0] > (1-k)^2 \prod_{t=2}^{\infty}\left(\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}}\right)^2 = \left((1-k)\prod_{t=2}^{\infty}\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}}\right)^2$. Now we show $\prod_{t=2}^{\infty}\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}} \geq \exp\left\{-\frac{k}{2k-1}\right\} > 0$ for $k > 1/2$.

$$\begin{aligned}
\prod_{t=2}^{\infty}\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}} &= \exp\left\{\log\left\{\prod_{t=2}^{\infty}\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}}\right\}\right\} = \exp\left\{\sum_{t=2}^{\infty}\log\left(\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}}\right)\right\} \\
&= \exp\left\{\sum_{t=2}^{\infty}\log\left(\left(\frac{k}{1-k}\right)^{t-2}\right) - \log\left(1+\left(\frac{k}{1-k}\right)^{t-2}\right)\right\} = \exp\left\{\sum_{t=2}^{\infty} -\frac{1}{z_t}\right\}
\end{aligned}$$

for $z_t \in \left(\left(\frac{k}{1-k}\right)^{t-2}, 1+\left(\frac{k}{1-k}\right)^{t-2}\right)$, by the Mean-Value Theorem. Hence

$$\prod_{t=2}^{\infty}\frac{\left(\frac{k}{1-k}\right)^{t-2}}{1+\left(\frac{k}{1-k}\right)^{t-2}} > \exp\left\{\sum_{t=0}^{\infty} -\frac{1}{\left(\frac{k}{1-k}\right)^{t}}\right\} = \exp\left\{-\frac{k}{2k-1}\right\}.$$

This establishes the lower bound on observing that the players forever play actions at least k despite the fact that they should be getting infinite evidence that the true state (and hence the optimal action) is 0.

While we prove the result in a model with a continuum of actions and a continuum of signals, it holds equally well in models with a finite number of actions and signals. (Of course, with finite actions, rational play can also be consistent with herding on the wrong state.)

Proposition 5 establishes that unlike rational players, INIT players may never come to learn the true state in our continuous model.[7] Yet it does not establish what INIT players eventually come to believe.

Claim: If any subsequence of beliefs converge, then they converge to 0 or 1.

Sketch of proof: Suppose beliefs converge to $p \in \left(\frac{1}{2}, 1\right)$. Fix ε small. If $a_{t-2} = a_{t-1} = p$, then $a_t > p + \varepsilon$ with high probability, so beliefs do not converge. If beliefs converge to $L^{-1}(1)$, then after $a_{t-1} = L^{-1}(1)$, $a_t = s_t$, and therefore beliefs do not converge.

Claim: Actions converge with probability one.

Proof: To be completed.

Beliefs and actions in INIT play converge faster than in Bayesian Nash equilibrium because of the disproportionate effect of the first few signals and actions. When beliefs and actions are slow to converge, it is likely to be because early signals were misleading. This suggests that persistent weak beliefs may be more likely wrong than right.

Claim: For each $\varepsilon > 0$, there exists $T \in \mathbb{N}$ such that if for each $t < T$, $a_t \in \left(\frac{1}{2} + \varepsilon, 1 - \varepsilon\right)$, then $\Pr[\omega = 0 | (a_1, \ldots, a_T)] > \Pr[\omega = 1 | (a_1, \ldots, a_T)]$.

Proof: Let $\varepsilon > 0$ be given. Define $S = \min \left\{ t : \frac{L^t\left(\frac{1}{2}+\varepsilon\right)}{1+L^t\left(\frac{1}{2}+\varepsilon\right)} \geq 1 - \varepsilon/2 \right\}$. Since $L\left(\frac{1}{2} + \varepsilon\right) > 1$ by the strict monotone likelihood ratio property, S exists. Now, if $a_{S+1} < 1 - \varepsilon$ in INIT play, then for some $\delta > 0$, $s_{S+1} < \frac{1}{2} - \delta$ (as $s_{S+1} \geq \frac{1}{2}$ leads to $a_{S+1} \geq 1 - \varepsilon/2$). Likewise, if $a_{S+1} \in \left(\frac{1}{2} + \varepsilon, 1 - \varepsilon\right)$ and $a_{S+2} < 1 - \varepsilon$, then $s_{S+2} < \frac{1}{2} - \delta$. Define $T = S + \min \left\{ t : \frac{L^S(1-\varepsilon)L^t\left(\frac{1}{2}-\delta\right)}{1+L^S(1-\varepsilon)L^t\left(\frac{1}{2}-\delta\right)} < 1/2 \right\}$, which

[7] The fact that rational play converges to the true state in this model can be proven along the lines of the proof by Goeree et al (2004) that beliefs converge to the true state in a quantal-response equilbrium in the binary-action model.

exists since $L^t\left(\frac{1}{2}-\delta\right) < 1$. If for each $t < T$, $a_t \in \left(\frac{1}{2}+\varepsilon, 1-\varepsilon\right)$, then $\Pr[\omega = 0] > \Pr[\omega = 1]$. as claimed.

The Proposition establishes that if players take actions above one-half for many periods—indicating that they all believe $\omega = 1$ more likely than $\omega = 0$—but no one takes an action close to one, then in fact it is more likely that $\omega = 0$. In INIT play, each player believes that all of her predecessors' actions coincide with their signals. A player at the end of a long run of high actions believes that her predecessors must all have high signals. So, the only reason why she would not conclude that $\omega = 1$ with virtual certainty must be that she receives a very low signal herself. Hence, the only times when players take actions above one-half for many periods without any of them taking an action sufficiently close to one are when, after a while, all of those players are receiving low signals, indicating that zero is more likely to be the state.

Table XXX illustrates this phenomenon for the case where $f_0(s) = 2(1-s)$ and $f_1(s) = 2s$. Simulating the ten-player cascade game XXX times for Bayesian Nash and INIT play, it shows the probabilities of various actions being played when the state is $\omega = 1$.

Player	$\Pr[a_t < 0.05]$ INIT	$\Pr[a_t < 0.05]$ BNE	$\Pr[a_t > 0.95]$ INIT	$\Pr[a_t < 0.95]$ BNE
1	0.0025	0.0026	0.0977	0.0976
2	0.0058	0.006	0.303	0.3035
3	0.0216	0.007	0.5965	0.4871
4	0.0483	0.0069	0.764	0.6247
5	0.0739	0.006	0.8332	0.7232
6	0.0914	0.0051	0.8623	0.7954
7	0.1016	0.0041	0.8754	0.8477
8	0.1068	0.0033	0.8815	0.8857
9	0.1098	0.0026	0.8845	0.9148
10	0.1115	0.002	0.8856	0.9356

Column 1 suggests that INIT play can lead to actions converging to zero when the true state is one, while Column 2 shows that this cannot occur with rational play. Summing entries in the last row from Columns 1 and 3 gives that 99.71% of INIT Player Tens play actions below 0.05 or above

0.95. By contrast, only 93.58% or rational Player Tens do. Although we have no formal result along these lines, the observation suggests that INIT play converges faster than rational play.

[INSERT TABLE ON SCHRAG EFFECT.]

Kübler and Weizsäcker (2004) experimentally explore a variant of the discrete cascade game with only two possible states and signals. They estimate the fraction of subjects who are level-k players in the sense of Stahl and Wilson (1995), where level-0 players uniformly mix over all available actions, and level-k players best respond to beliefs that all other players are level $k-1$.

Proposition: INIT play in the cascade game coincides with the prediction of Stahl and Wilson's level-two types.

Proof: A level-2 type best responds to beliefs that all other players best respond to all players' mixing uniformly on $[0,1]$. Hence, level-1 types believe that predecessors' actions convey no information about their private information and therefore act in the same way as fully-cursed agents. Level-2 types then best respond to fully-cursed behavior on the part of all other players, the definition of INIT non-equilibrium.

Kübler and Weizsäcker's (2004) experiment follows a model of Anderson and Holt (1997) that differs from our baseline model in two ways: first, there are only two possible states of the world and signals; second, each player must pay a small cost in order to receive her signal, and those choosing not to pay this cost do not receive signals. Player k can observe all of her predecessors choice of urn but none of their choice whether to receive signals. It is straightforward to show that in an INIT equilibrium, Player 1 buys a signal while all later players do not buy signals but mistakenly believe that all of their predecessors have bought signals.[8]

[8] What is the INIT non-equil in KW? Player 1 pays the cost to acquire a signal, which she follows. Player 2 thinks Player 1 is fully cursed, in which case Player 1 would behave exactly as she does. When deciding whether to acquire a signal himself, Player 2 recognises that if his signal agreed with Player 1's action, then he would prefer to follow Player 1's action; if his signal disagrees with Player 1's action, then he would be indifferent between following Player 1's action and following his signal. Since following Player 1 turns out to be optimal regardless of his signal, Player 2 does not purchase a signal. Now, if Player 3 is an INIT Player, he believes Player 2 to be fully cursed, in which case Player 2 would purchase a signal and follow it. A fully-cursed Player 2 would

Like in the baseline model, INIT non-equilibrium coincides with the predicted play when all players are level-2 players in the sense of Stahl and Wilson (1995). Rather than estimate the distribution of these levels directly, Kübler and Weizsäcker (2004) estimate a variant of Quantal Response equilibrium that relaxes the assumption that players have rational expectations about the distribution of other players' actions (see Weizsäcker, 2003); instead, Player k may believe that Player $k-1$ best responds with different noise than he does to Player $k-1$ private information and observation of his predecessors. What they find is that players believe that their predecessors use noisier best replies than they do themselves, consistent with the notion that players believe their predecessors to be more cursed than they themselves are.

5 Conclusion

To be written.