

The Timing of Social Learning

Brian W. Rogers¹

First Draft: May, 2003

Revised: September, 2005

Summary. We consider a pure informational externalities environment in which agents make binary decisions. The agents are asymmetrically informed about a payoff-relevant state variable, and choose the timing of their decisions strategically. There is a delay cost that must be balanced against the possibility of learning from the announcements of predecessors. For two player games in discrete time, we show that in the unique equilibrium the game ends in finite time and the agent with better information decides earlier. As the time intervals become vanishingly short, all announcements occur immediately, no delay costs are incurred, and the equilibrium outcome approaches the first best. We show that for games with many players and short time intervals, the true state is revealed immediately and thus an efficient herd arises in which almost all agents announce correctly. For any time interval, welfare is higher under strategic timing relative to an exogenous sequence of decisions.

Keywords and Phrases: Social learning, Strategic timing, Information aggregation, Information cascades, Herding

JEL Classification Numbers:

¹Division of Humanities and Social Sciences, California Institute of Technology, Mail code 228-77, Pasadena, CA 91125, USA (email: rogers@hss.caltech.edu, phone: (626) 395-4086, fax: (626) 405-9841). I am grateful to John Duggan, Matthew Jackson, Thomas Palfrey, and especially Jacob Goeree for guidance and suggestions. Elena Asparouhova, Jernej Copic, P.J. Healy and Betsy Sinclair also provided useful comments.

1 Introduction

When individuals share a common objective but are faced with limited information, they may base their decisions partly on the observable actions of others. For instance, portfolio managers tend to bias their investment decisions towards the decisions of their colleagues.¹ In seminal papers, Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) explore settings where each of a set of ex-ante identical agents receive private signals and publicly guess the outcome of a state variable in a predetermined sequence. The common prediction of these models is that at some point a predominance of choices for a particular action overwhelms any privately held information. At this point all remaining decision makers also take the same action, even if it is contradicted by their signals.

It is this phenomenon that has come to be known as an *information cascade*. A cascade entails an informational inefficiency in that once decisions do not reflect signals, there is no mechanism by which private information can be distributed to other agents and aggregated, and the process of social learning halts. Thus even in large economies where there is an abundance of distributed information, with substantial probability most individuals make suboptimal decisions.

In this standard setup, the order in which agents make decisions is fixed exogenously. Yet there may be interesting phenomena that drive the observed timing of decisions in the field. One factor is that as the number of announcements grows, inferences based on these decisions become more accurate on average. Consequently, individuals who decide later in the sequence expect to do better.² On the other hand, there is typically an explicit cost to delay. For instance, a financial manager making an investment decision pays the cost of holding a suboptimal portfolio while collecting the information to make her decision.³ As another example, consider the decision by a group of firms about whether to adopt a new production technology.⁴ The profitability of the technology is uncertain. Each firm would like to observe the adoption by other firms before making their decision, as this provides valuable information. Yet delaying the decision is costly since there is a possibility that the firm is using a suboptimal technology in the

¹See, for example, Scharfstein and Stein (1990).

²Gale (1996) explains this logic in some detail, and also provides an overview of related phenomena from the herding literature.

³Dasgupta (2000) examines this setting in more detail.

⁴See Kapur (1994) for an analysis of technology diffusion.

interim. Thus there is an inherent tension between deciding early and late that must be resolved. At equilibrium, agents balance this tradeoff optimally, realizing that others are simultaneously making the same calculations. This interaction is modeled below by allowing timing decisions to be strategic variables in a social learning game.

The model presented below may be summarized as follows. Agents are asymmetrically informed about a binary state variable. Signals are conditionally independent and take values on the unit interval, with a monotonic likelihood ratio. Agents must publicly guess the outcome of the state variable at a time of their choosing. The payoff is positive only when the guess is correct, and future payoffs are discounted at a common rate so that waiting is costly. We characterize equilibrium outcomes in this setting and the implications for efficiency.

Notice that, as opposed to the canonical model, the accuracy of private information is heterogeneous. Signals vary in quality, ranging from uninformative to fully revealing the state variable.⁵ The combination of strategic timing and differential signal quality drive the intuition of the analysis below. To illustrate, consider the situation of a worker seeking employment, where job offers are publicly observable. It seems likely that some potential hirers would have better information regarding the worker's quality than others. For instance, some firms may gain more information from the interview, or may personally know those who wrote recommendation letters. Further, firms with better information could use this difference to gain a strategic advantage over other firms in the hiring process, e.g. by making quicker decisions. Another example derives from bank runs, as experienced in the United States before the institution of federal deposit insurance.⁶ Under this interpretation of the model, the timing of deposit withdrawals is related to an individual's private information. Those with the strongest beliefs that the bank is insolvent should be the first to act. But taking this action is costly if the bank does not fail, so there is an incentive to delay in order to observe the others' decisions.

This paper provides an explanation of how rational agents with differing information qualities behave when they strategically choose the timing of their decisions. Since utility depends only on an agent's own action and the state of the world, all strategic interactions are due solely to an informational externality. We show that in two player games there is a unique equilibrium. It

⁵Thus signals are unbounded in the sense of Milgrom (1979). Smith and Sorenson (2000) contains a thorough analysis of social learning in the case of unbounded signals.

⁶Caplin and Leahy (1994) studies a model of market dynamics with an application to bank runs.

exhibits monotonic “sorting”, i.e., agents with better information decide earlier, and those with lower quality signals behave more patiently and incur more delay cost in order to learn from others. In this model, continuously varying signal qualities are important for the existence of a symmetric equilibrium in pure strategies, since signal quality differentiates agents and breaks indifferences about the timing decision that would otherwise occur. As the time grid becomes increasingly fine, the equilibria approach a limit where all action occurs immediately, and delay costs vanish. Since the sorting improves as the time intervals shrink, the agent with the lower quality signal learns from the better informed agent with probability converging to one, and the equilibrium outcome approaches the first best. However, this limit outcome can only be supported as an equilibrium in continuous time under a particular “tie-breaking” assumption regarding simultaneous decisions. Further, since the equilibria in the discrete games are unique, there is no other approachable equilibrium in the continuous game.⁷

The main emphasis of this paper concerns how the welfare properties of social learning are affected by the strategic timing of decisions relative to exogenous sequencing, for which welfare implications are very negative. Despite the equilibrium phenomenon of herding, information is aggregated efficiently in the present model due to the sorting property when there are only two agents. Yet these outcomes are not fully efficient because of the delay cost implicit in the sorting. We find that efficiency increases monotonically as the time intervals shrink, and, as mentioned above, approaches the first best outcome in the limit. For any finite time interval, aggregate welfare is higher in the endogenous timing model than in the standard exogenous sequence model. In games with many players full efficiency is approximated when the time intervals are short. The reason is that with unbounded signals, there are always some signals that produce a dominant strategy to announce immediately, and with many players some individuals receive such signals and reveal the true state to others with high probability. Thus the message of this paper is positive in the sense that we view the assumptions of this model as more appropriate in many settings, and welfare is unambiguously higher in this setting.

The rest of the paper is organized as follows. The next section describes in more detail how the present model relates to existing literature, and highlights this paper’s contributions. Section 3

⁷Several authors have provided general existence theorems for discontinuous games, including (Reny (1999), Jackson, Simon, Swinkels, Zame (2002), Dasgupta and Maskin (1986a,b) and Milgrom and Weber (1985)). Despite the fact that none of these results are readily applicable to the present model, we are also able to find asymmetric equilibria, although they are not perfect.

presents the model for the case of two players. Section 4 contains the equilibrium characterization and the convergence result that as time intervals vanish the game ends immediately. Section 5 then generalizes the model to allow for a large number of players, and analyzes equilibrium outcomes in that context. The comparison of information aggregation and efficiency to the case of exogenous ordering of decisions is handled in Section 6. Some extensions and limitations of the model are discussed in Section 7. Finally, section 8 concludes. Proofs of some main results and technical details are included in the Appendix.

2 Related Literature

A number of authors, including Chamley and Gale (1994), Gul and Lundholm (1995), Chamley (2004), and Zhang (1997), have also analyzed models of social learning in which timing decisions are strategic variables. Interestingly, the results of these models are mixed, and in contrast to the exogenous sequencing models, depend crucially on modeling assumptions about the nature of time, uncertainty, and the richness of the type and action spaces. Consequently, it is difficult to draw unambiguous conclusions regarding welfare implications.

Chamley and Gale (1994) consider an investment model where there is only one action, but infinite delay is not a dominated strategy. The most important difference from the current paper is that there is no notion of differentiated signal quality. Because of this, social learning does not lead to welfare gains in their model. The reason is that the only symmetric equilibrium involves mixing probabilities that cause agents to be indifferent between announcing and delaying, so that informational gains are exactly balanced by delay costs in equilibrium. One reason for the welfare improvement in the present paper is that the timing of announcements reveals information about signal qualities, which can not happen in Chamley and Gale (1994). A common finding of these models, however, is that the game ends quickly as the time interval shrinks.

The model of Gul and Lundholm (1995) operates in continuous time and uses a continuous action space, and again signals are not differentiated by quality. The focus of that study is on “clustering”, the tendency of announcements to be closer together than if they were made independently. Under their utility specification they find that social welfare is lower than in the standard exogenous sequencing model, which is the opposite result of the present paper. As before, time can not serve as a screening device for worse informed agents to learn. A second

effect is that, since time is continuous, the relevant comparison to the exogenous sequencing case involves no discounting for later announcers. This suggests one sense in which the modeling of time as a discrete quantity may be more appropriate in social learning models.

Building on the work of Chamley and Gale (1994), Chamley (2004) incorporates differential signal quality into the model, which generates results more similar to this model: (i) as the time intervals vanish, the game ends immediately, and (ii) the true state is asymptotically revealed as the number of players diverges. Interestingly, though, he finds a multiplicity of equilibria with widely varying information aggregation properties. This contrasts with the present model for which we find a unique equilibrium⁸ with unambiguous welfare implications. For instance, a rush in which all agents announce immediately constitutes an equilibrium in Chamley (2004) but not in the present paper. This contrast is due to the fact that there are two possible announcements in the present context, so that agents with neutral beliefs will find it optimal to delay given that others are announcing. Another difference resulting from the number of possible announcements is that if ever there is a period in which no agent makes an announcement, the game effectively ends.⁹

The most closely related work to the present model is Zhang (1996). The signal structure there is more general than the one defined below, with the important difference that signals there are bounded, whereas in this model there is a possibility of being arbitrarily well informed. Action takes place in continuous time. In the unique symmetric equilibrium there is delay, followed by an immediate cascade based on the (single) highest quality signal. This result, however, depends crucially on an additional source of uncertainty that has no counterpart here. The present model also has a unique equilibrium, without restricting attention to symmetric strategy profiles, which does not converge to the continuous time result found by Zhang. Further, with many players herding is inefficient in Zhang’s model whereas the opposite is true here. These differences relate to the presence of unbounded signals in my model. With many players, unbounded signals guarantees that some agents announce early on, and the information conveyed

⁸This is without restricting attention to symmetric equilibria a priori, as do Chamley and Gale (1994) and Gul and Lundholm (1995).

⁹This property is due to the fact that there is only one kind of announcement, and so a period with no announcements reveals the worst possible news about the profitability of announcing. As such, the equilibrium analysis in Chamley (2004) can make use of a “two-step” property which effectively allows the payoff from announcing to be compared with the payoff from delaying exactly one period, rather than having to compute the continuation value of delay as a function of others’ strategies over all future periods.

from their choices reveals the state variable with high probability. As a result, the asymptotic welfare properties are opposite in the two models. Despite the similarity of these models, Chamley (2004) argues for the importance of analyzing a discrete time version of Zhang's model.¹⁰ In fact, Chamley (2004) conjectures that there may be multiple equilibria in such a game. This paper suggests that this is not true.

One contribution of this paper is that, when combined with other results from the previous literature, one now obtains a rather complete picture of how various modeling assumptions about information and timing translate into different results. With respect to the information structure, two general points can be made. First, when signals are unbounded, the implications for welfare in large economies are positive. With arbitrarily accurate signals and many individuals, there will always be some individuals who are sufficiently sure about the true state that it can not benefit them to delay, and their decisions collectively reveal the state to all other individuals, who then have nothing else to gain and so announce correctly immediately following the first announcers. In contrast, when signals are bounded the welfare results remain negative: inefficient herds may still arise under strategic timing since the option value of delay must balance the cost. Second, when signal quality is homogenous, symmetric equilibria typically involve mixed strategies in order to balance the marginal benefits and costs of delay to maintain indifference, whereas with a continuum of signal qualities the model may admit symmetric pure strategy equilibria where agents with better qualities announce earlier.

With respect to the modeling of time, one finding is that equilibria in discrete time games may not converge to the equilibria of corresponding continuous time games. For instance, a result that is somewhat robust in the literature is that as discrete time intervals vanish, all action takes place instantly in the limit, whereas such equilibria are not easily supported in continuous time. Discrete time has the advantage that it is easier to interpret, in the sense that how information is processed and responded to is more readily defined.

¹⁰See footnote 13 in that paper.

3 Model

There is a countable set of dates $T^r = \{0, r, 2r, \dots\} \cup \{\infty\}$, with the interval between consecutive dates given by $r > 0$, at which agents in a set $N = \{1, 2\}$ make decisions.¹¹ The payoff-relevant states of the world are given by $\omega \in \Omega = \{A, B\}$. Agents have uniform prior beliefs regarding the true state, so that $\Pr(A) = \Pr(B) = 1/2$. The model is one of incomplete information. Agents receive signals $p_i \in [0, 1]$, $i = 1, 2$.¹² Conditional on ω , the p_i 's are independently and identically distributed with *pdf*

$$f(p|\omega) = \begin{cases} f(p|A) = 2p & , \text{ if } \omega = A \\ f(p|B) = 2(1-p) & , \text{ if } \omega = B \end{cases} \text{ with support } [0, 1].$$

Thus low realizations of p are evidence of state B while high realizations are evidence of A . In particular, a signal of $p = 1/2$ is uninformative, while signals of $p = 0$ or $p = 1$ perfectly reveal the true state, so that the signals are unbounded.¹³ The unconditional distribution of types is given by

$$f(p) = \frac{1}{2} (f(p|A) + f(p|B)).$$

which is uniform on the unit interval. The most important feature of the signal structure is the monotone likelihood ratio. More general structures are discussed in Section 8. All elements of the model, except the realizations of the signals and the state variable, are common knowledge among the agents.

Each agent must guess the state of the world, and so chooses a $c_i \in C = \{a, b\}$ and a time $t_i \in T$ at which to make her announcement. Denote by $x_i^t \in X \equiv C \cup \{0\}$ agent i 's action at date t , with the interpretation that $x_i^t = 0$ records the event that i did not make an announcement at t , and let $x^t = (x_1^t, x_2^t)$ be the profile of actions at time t . A history h^t at time t specifies whether or not each agent has made an announcement, and if so, what the announcement was and the time at which it was made. Thus we can write $h^t = \cup_{\tau \in T: \tau < t} x^\tau$, with the convention that $h^0 = \emptyset$.¹⁴ A strategy for i specifies an action to take at every history. Let $\mathcal{P}(Z)$ denote

¹¹The notation T is used in place of T^r when no confusion arises.

¹²We use the phrases “agent with signal p ” and “agent of type p ” interchangeably for convenience, even though the language “signal” is more accurate.

¹³Smith and Sorenson (2000) analyze social learning in a general framework with unbounded signals and exogenous ordering of decisions.

¹⁴To formally express the fact that each agent can make only one announcement, write \mathcal{H}^t for the set of

the set of probability distributions over a set Z . A *behavioral strategy* σ_i for i maps signals and histories into probability distributions over X , specifying the probabilities that i announces a , announces b , and delays, respectively.¹⁵ Denote by $\sigma = (\sigma_1, \sigma_2)$ the profile of strategies.

Agents receive a payoff of one if they guess correctly and zero otherwise. A common instantaneous discount factor $\gamma > 0$ accounts for temporal preferences, so that the present value of guessing correctly at time t is $\exp(-\gamma t)$. Denote the game so defined by G^r .

Given the instantaneous discount rate γ and interval length r , one immediately infers that $\delta \equiv \exp(-\gamma r)$ is the effective period-to-period discount factor. Treating temporal preferences as fixed and exogenous, there is a one-to-one correspondence between δ and r , such that $\delta \rightarrow 1$ as $r \rightarrow 0$. In discussing the results below, we make use of this equivalence.

Agents use Bayes' Law to update their beliefs. Let $\pi_i(p_i, t, h^t, \sigma)$ represent i 's belief that the true state is A at time t conditional on i 's type p_i and on the history h^t , given the strategy profile σ . Thus after observing her type and before the game starts, agent i 's posterior belief is given by

$$\pi_i(p_i, 0, \emptyset, \sigma) = \frac{f(p_i|A)P(A)}{f(p_i|A)P(A) + f(p_i|B)P(B)} = p_i.$$

We use perfect Bayes Nash equilibrium (PBNE) as the solution concept.¹⁶ In the next section we find a unique PBNE. In fact, the game is dominance solvable, so that the equilibrium strategy profile is determined through iterative elimination of strictly dominated strategies.

Although the set of strategies is quite large, a few observations greatly restrict the set of possible equilibrium strategies. First, note that if the other agent announces first, the unique best response is to make a decision immediately thereafter since there is no possible benefit in delaying further, and delay is costly. Thus after one agent announces, the other agent announces as quickly as possible, payoffs are realized, and the game effectively ends. Second, note that if an agent announces before her opponent, she should always announce a if and only if $\pi_i(p_i, t, h^t, \sigma) \geq$

possible histories at time t and let $\mathcal{H} = \cup_{t \in T} \mathcal{H}^t$ denote the space of all histories. We require that for all $h^t \in \mathcal{H}$, $x_i^\tau \neq 0 \Rightarrow x_i^{\tau'} = 0$ for all $\tau < \tau' < t$.

¹⁵The restriction that only one announcement may be made can be written as follows: for all $p_i \in [0, 1]$ and all $h^t \in \mathcal{H}$, $\sigma_i(p_i, h^t) = (0, 0, 1)$ whenever there exists $\tau < t$ such that $x_i^\tau \neq 0$.

¹⁶Let $g_i(\cdot | p_i, t, h^t, \sigma)$ denote i 's probability assessment of j 's type given her own type, the time, the current history, and strategy profile σ . A PBNE is characterized by a strategy profile σ and a belief profile g such that for $i = 1, 2$, $j \neq i$, (i) σ_i specifies a best response to σ_j at every history given beliefs g_i , and (ii) g_i is consistent with Bayes' Law at every history that is reached with positive probability under σ .

$1/2$.¹⁷ Given these two simple observations, the equilibrium strategies for i can then be fully described by a mapping $s_i : [0, 1] \rightarrow \mathcal{P}(T)$ that gives the distribution of stopping times at which an agent will announce as a function of her type p_i provided the other agent has not yet announced. We will be primarily concerned with pure strategies, which in this reduced form may be expressed by a measurable function $s_i : [0, 1] \rightarrow T$. That is, i is willing to wait only until time $s_i(p_i)$ to make her announcement, and will announce earlier if and only if her opponent announces first. Clearly, any equilibrium strategy must satisfy $s_i(0) = s_i(1) = 0$, as an agent who knows the true state with certainty has nothing to gain by delaying, and would still incur the cost, i.e., she has a strictly dominant strategy.

In what follows we will be predominantly concerned with symmetric equilibria. In such cases, subscripts on strategies and beliefs are dropped when no confusion arises. In addition, the equilibrium strategies we find treat the two states of the world symmetrically. Formally, we have the following

Definition 1 *A strategy s_i for i is information-symmetric if for all $0 \leq p \leq 1/2$, $s_i(p) = s_i(1 - p)$.*

Thus information-symmetric strategies are symmetric mappings about the uninformative signal ($p = 1/2$), which means that we can restrict attention to the *quality* of an agent's type, as measured by $q = |p - 1/2|$. Note that when agents use information-symmetric strategies, beliefs about the true state $\pi_i(p_i, t, h^t, s)$ change only when an announcement is made. To see this, note that whenever a type $p < 1/2$ plans to announce at a time t , there is a corresponding type $1 - p$ who also plans to announce at t . Thus agents learn nothing about the likelihood of the true state through the passage of time, unless an announcement is made.¹⁸

We refer to PBNE that are both symmetric and information-symmetric as SISPBNE. Since under information-symmetric strategies we can restrict attention to types $p \leq 1/2$, a SISPBNE in pure strategies (pure SISPBNE) is fully described by a function $s^* : [0, 1/2] \rightarrow T$. So s^* gives the planned stopping time for each signal quality an agent may have.

¹⁷The results are not sensitive to how we specify an agent's decision when she is indifferent between the two choices, as this occurs with zero probability.

¹⁸See the Appendix for a formal statement and proof of this important fact.

4 Equilibrium Analysis

The equilibrium strategies below satisfy the following

Definition 2 *Given $\gamma > 0$, fix $r > 0$ and an information-symmetric pure strategy s_i for i in G^r . s_i is characterized by cutpoints if there exist a number $k < \infty$ and numbers $0 = p_0 < p_1 < \dots < p_k = 1/2$ such that if $p_{l-1} < p < p_l$ then $s_i(p) = (l - 1)r$, $l = 1, \dots, k$.*

Notice that any strategy characterized by cutpoints satisfies two properties. First, there can be no “empty dates”, meaning that for any $l < k$, there exists a set of positive measure of types for which $s(p) = lr$. Second, a monotonicity condition is satisfied, in the sense that as an agent’s signal quality q increases, her stopping time (weakly) decreases. We can now characterize the equilibria of G^r . The following proposition shows that G^r has a unique equilibrium, and that it is symmetric, information-symmetric, in pure strategies, and characterized by cutpoints.

Proposition 1 *Fix $r > 0$ and $\gamma > 0$. G^r has an (essentially) unique PBNE s^* . Moreover, s^* is symmetric, information-symmetric, in pure strategies, and characterized by cutpoints.*

Proof. See Appendix. ■

The intuition for Proposition 1 is as follows. That there exists a symmetric and information-symmetric equilibrium results from the ex-ante symmetry of the game. Uniqueness is obtained by iteratively eliminating strictly dominated strategies in the following way. At time zero, there is a group of types of each agent who are sufficiently well informed that it can not possibly benefit them to wait – the cost of delaying is too high even if by waiting one period they could learn the outcome of the state variable. Given that these types must announce at time zero in any equilibrium, it turns out there is enough expected benefit from waiting to induce all less well informed agents to delay. This is shown through iteratively eliminating the strategies “announce immediately with positive probability” for types who are increasingly well informed, starting from the uninformative type ($p = 1/2$), until only the types who had the dominant strategy to announce remain. Then at the next date, if neither agent has announced, it is common knowledge that both agents are not so well informed that they dropped out at the previous date. The same logic is applied again, so that of those types remaining, some agents

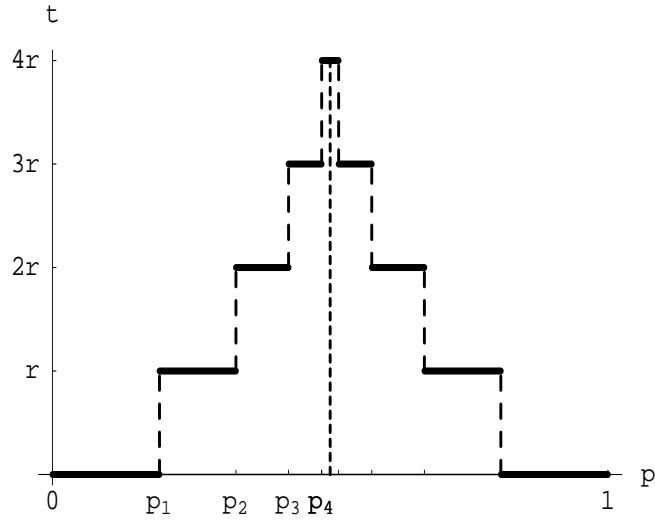


Figure 1: The equilibrium strategy $s^*(p)$ as a function of type p for $r = 0.045$ and $\gamma = 1$.

can not possibly gain enough from waiting to offset the cost, etc. This construction results in a unique equilibrium. Figure 1 depicts a typical equilibrium strategy. In this case, the maximum length of the game in equilibrium is $4r$.

Proposition 1 provides a strong result characterizing the equilibrium of G^r . Note that for any fixed temporal preference γ , as r decreases, the cost of delaying one period decreases, and so agents become more willing to delay on the margin. This means that as r tends to zero, the cutpoints depicted in Figure 1 shift continuously towards the boundaries of the unit interval, additional cutpoints enter from the midpoint $p = 1/2$, and the number of potential periods grows without bound. Since the marginal cost of delaying one period approaches zero as time progresses, one could imagine that for sufficiently small r , the game could last an infinite number of periods with positive probability. The next result shows that this is in fact not possible, i.e., that G^r must end in finite time for every positive r .

Proposition 2 *For every $\gamma > 0$ and $r > 0$, G^r ends in finite time. That is, there exists $k < \infty$ such that $s^*(p) < kr$ for all $p \in [0, 1]$ in the unique equilibrium s^* .*

Proof. See Appendix. ■

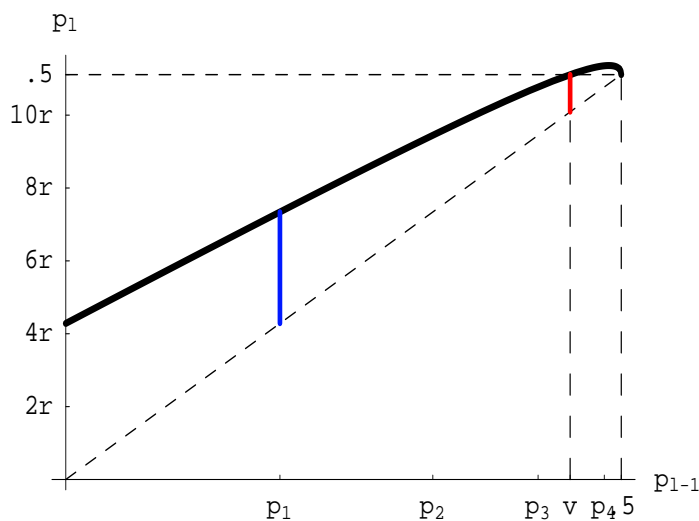


Figure 2: The equilibrium cutpoint function $p_l(p_{l-1}; 0.956)$.

The proof is established as follows. First from the proof of Proposition 1 (see Appendix), we have an equation that implicitly defines the cutpoint p_l as a function of the previous cutpoint p_{l-1} for any $\delta < 1$. This equation may be solved to provide an explicit function $p_l = p_l(p_{l-1}; \delta)$, which is graphed in Figure 2. The properties of this function can be used to derive a positive minimum distance between any consecutive cutpoints p_l and p_{l-1} , i.e., a minimum interval length $d(\delta)$ of types who plan to announce at time $(l-1)r$. Graphically, the interval length $p_l(p_{l-1}; \delta) - p_{l-1}$ is the vertical distance between the bold curve $p_l(p_{l-1}; \delta)$ and the dashed 45° line at the point $p = p_{l-1}$. The cutpoint function takes valid values $p_l \leq 1/2$ up to the point $p_{l-1} = v$, where the function crosses the value $\frac{1}{2}$. Thus the red line segment depicts the smallest possible interval length d for a given δ , which is always positive. Given this bound, the maximal number of equilibrium intervals, which corresponds to the maximum number of periods the game can last, is $1/2d$.

One consequence of Proposition 1 is that unless the agents happen to announce simultaneously, they will necessarily make the same decision. To see why this is so it is enough to consider the decision problem facing an agent whose opponent has already announced. Given the sorting nature of equilibrium, the first announcer has strictly higher signal quality, so the posterior beliefs of the second announcer are dominated by the information from the decision

of the first announcer, independent of which state the second announcer's private information indicates. That is, in equilibrium, $\pi_i(p_i, t, h^t, s^*) \geq 1/2$ if $x_j^\tau = a$ for some $\tau < t$. Thus the second announcer finds it optimal to make the same announcement as the first announcer. The only way in which the announcements of the two agents can differ is if they decide simultaneously, in which case they have approximately the same quality of information, and neither agent learns from the announcement of the other.

Having characterized the equilibrium, we now turn attention to analyzing the comparative statics as r tends to zero. By the results of the previous section, we know that the equilibrium strategy s^* partitions the set of types $p \leq 1/2$ into intervals, where each interval of types is willing to wait until the same time to announce, conditional on the other agent having not previously announced. For a fixed δ , call the number of equilibrium cutpoints $k(\delta)$. Thus the game necessarily ends by time $rk(\delta)$, since at each date there must be some interval of types planning to announce. By Proposition 2, $k(\delta) < \infty$ for all $\delta < 1$. Note however that $k(\delta)$ diverges to infinity as δ approaches unity, meaning that as the time intervals become vanishingly small, the number of periods that the game potentially lasts becomes infinite. The relative rates at which r and k change determines the limiting behavior of equilibrium timing.

The next result shows that as r becomes vanishingly small, the game ends immediately.

Theorem 1 *For every $\gamma > 0$, G^r ends immediately as the time interval vanishes. That is, $\lim_{r \rightarrow 0} rk(\delta) = 0$.*

Proof. See Appendix. ■

Theorem 1 is established by deriving an upper bound $\bar{k}(\delta)$ on the number of equilibrium cutpoints, and showing that $r\bar{k}(\delta)$ limits to zero. The proof is essentially an application of L'Hopital's rule; it is constructing the upper bound $\bar{k}(\delta)$ that requires work. As in the proof of Proposition 2, the properties of the function $p_l(p_{l-1}; \delta)$ (see (9.4) in the Appendix) are exploited. Using a linear approximation to $p_l(p_{l-1}; \delta)$ instead of the actual function, one can define alternative equilibrium cutpoints z_l with the property that $z_l \leq p_l$ for all $l \leq k$. Because of the linearity, the recurrence relation that defines the $\{z_l\}$ can be solved. Setting this expression equal to $1/2$ defines $\bar{k}(\delta)$, the number of cutpoints that result by using the approximation to $p_l(p_{l-1}; \delta)$, which is an upper bound on the actual number of cutpoints.

5 Many Players

In this section we expand the model to allow for an arbitrary finite number of agents. Let the set of agents be $N = \{1, 2, \dots, n\}$, $n < \infty$. Denote these games by G_n^r , where all other aspects of the game are unchanged. The details of this extended model are difficult to work out precisely, so we are not able to explicitly characterize the equilibria as in the case of two players. The main difficulty is that with many players, there are many subgames that must be considered at all dates after $t = 0$. This problem is avoided when $n = 2$ because there is only one interesting history to consider at any date, namely, that both players have not yet announced. In contrast, there are many nondegenerate histories with several players.

When there are more players, there is more information that may be learned from the announcements of others. Thus one would expect individuals to behave more patiently, that is, to wait longer in equilibrium, all else equal. While this is in fact true, we can show that as the number of players grows large, the game necessarily ends quickly. The intuition is that no matter how many players there are, there is always a positive mass of types with good information who cannot be induced to delay $t = 0$, since even if they could learn the true state with certainty by delaying one period, the cost of delay (given $r > 0$) outweighs this benefit. As n grows, the probability of at least one player having such a signal approaches one. Thus as the number of players becomes large, the announcements of these players reveal the true state with near certainty after the first date $t = 0$. Let $L(\mathbf{p})$ denote the maximum realized length of G_n^r in any equilibrium given the n -vector of signals $\mathbf{p} = (p_1, p_2, \dots, p_n)$. Our next result states that for any positive r , the probability that the game ends after one period converges to one as n diverges.

Proposition 3 *Fix $r > 0$ and $\gamma > 0$. For every $\epsilon > 0$ there exists \bar{n} such that $\Pr(L(\mathbf{p}) = r) > 1 - \epsilon$ for all $n > \bar{n}$.*

Proof. By definition of G_n^r , the maximal continuation value of delay at $t = 0$ is $\delta = \exp(-\gamma r) < 1$. Thus for all $p < 1 - \delta$ and $p > \delta$ announcing at 0 is a strictly dominant strategy for every n . Thus for each n there will be some cutpoint $1 - \delta \leq p_1^n \leq 1/2$ such that all $p < p_1^n$ announce at $t = 0$. Assume that $\omega = B$ (the case of $\omega = A$ is parallel). Note that $\beta \equiv \Pr(p_i < p_1^n | \omega = B) > \Pr(p_i > 1 - p_1^n | \omega = B)$, since in state B low signals are more likely. Let $n_B = \#\{i | p_i < p_1^n\}$ be the number of agents who announce b at 0 as a function of \mathbf{p} . By a

law of large numbers, for every $\epsilon > 0$, $\lim_{n \rightarrow \infty} \Pr(|n_b/n - \beta| < \epsilon) = 1$. It follows that for every $\epsilon > 0$, $\lim_{n \rightarrow \infty} \Pr(\tilde{\pi}_i(p_i, x^0) < \epsilon) = 1$, where $\tilde{\pi}_i(p_i, x^0)$ represents the posterior obtained from signal p_i and actions x^0 . Thus as $n \rightarrow \infty$ beliefs converge to the true state after time 0, and so the marginal benefit of delay at the next date r approaches 0 since there is nothing more to be learned. Therefore all remaining types find it optimal to announce at $t = r$, since the cost of delay is strictly positive. ■

Since Proposition 3 is true for all $r > 0$, we can also state as a corollary that as r goes to zero, the probability that the game ends almost immediately is arbitrarily close to 1 for sufficiently large n . Note also that herding is observed in the equilibrium outcomes of G_n^r for large n . In particular, all agents who announce at time 0 do so independently from each other (so that most but not all of these agents announce correctly), but then all agents who delayed at 0 learn the true state and make the same optimal decision at time r . In contrast to the standard cascade model, where the order of announcements is fixed, there is no informational inefficiency associated with the herding. In fact, the opposite is true: all agents in the herd are announcing correctly with arbitrarily high probability. Observe also that as r tends to zero, the proportion of agents who decide at $t = 0$ vanishes, and these agents announce correctly with probability converging to 1. Thus full efficiency is approximated in large economies for small r . These observations are summarized in the following

Proposition 4 *Fix $\gamma > 0$. For every $\bar{t} > 0$, $\epsilon > 0$, and $\eta < 1$, there exists an $r > 0$ and \bar{n} that satisfy*

(i) *delay efficiency: $\Pr(L(\mathbf{p}) < t) > 1 - \epsilon$, and*

(ii) *informational efficiency: $\Pr(\frac{\#\{i | x_i = \omega\}}{n} > \eta) > 1 - \epsilon$ for all $n > \bar{n}$.*

Proof. By Proposition 3 for any $r < \bar{t}$ there exists \bar{n} to satisfy the first claim. For any r , the first cutpoint p_1 approaches $\delta = \exp(-\gamma r)$ as n diverges. Therefore, we can choose r small enough and \bar{n} big enough so that the individuals announcing at time zero (i) are an arbitrarily small proportion of individuals, and (ii) are arbitrarily well informed. This implies that an arbitrarily high proportion of individuals announce correctly with arbitrarily high probability, satisfying the second claim. ■

Given γ , for any n there is a maximum amount of time \bar{r}_n that a player with an uninformative signal is willing to wait to see the others' announcements. When $r = \bar{r}_n$ agents with every other signal strictly prefer to announce immediately, rather than waiting until time r . When $n = 2$, an uninformed player compares the expected utility of announcing at zero, which is $1/2$, to the expected utility of waiting, in which case he sees the other player's announcement, resulting in an expected utility of $(3/4)\delta$, so that $\bar{r}_2 = (1/\gamma) \ln(3/2)$. By the argument above, as n grows, the expected utility of delay when $r = \bar{r}_n$ converges up to δ , so that \bar{r}_n approaches $(1/\gamma) \ln(2)$. It is in this sense that players become more patient as n increases. However since the amount of time agents are willing to wait does not increase too dramatically with n , it may be possible to bound the length of G_n^r for arbitrary values of n , although we do not have this result.

6 Information Aggregation and Efficiency

The model aims to provide an explanation of how rational agents with heterogeneous signal qualities will use time as a screening device to learn about others' private information. We have determined that when agents choose the timing of their decisions strategically, they order themselves in a manner such that those with the best information decide first, and those with worse information are willing to wait in order to learn from early decision makers. Thus it is natural to measure the extent to which information is aggregated more efficiently in the endogenous timing model relative to the standard social learning setup, where an exogenous order is assigned to the players.

For tractability we consider the limiting case $\delta \rightarrow 1$ with two players. We first examine the welfare properties of the standard exogenous sequencing model as a benchmark for comparison to the present model. To do so, we suppose that before the game starts, and independently of the signal realizations, agents are randomly assigned a position in the queue, so that each agent is equally likely to be the first or second announcer.

Let Cor_i denote the event that i announces the correct state, i.e., $x_i = \omega$. The probability that the first agent guesses correctly is given by

$$\Pr(Cor_1) = \Pr(Cor_1|A) = \Pr(p_1 \geq 1/2|A) = \int_{1/2}^1 f(p|A)dp = 3/4.$$

Given the optimal decision rule of player 1, the probability that player 2 announces correctly is $13/16$.¹⁹ Note that this is consistent with the observation made above that agents who decide later in the sequence do better on average, in the sense that they announce correctly with higher probability. But deciding later is also costly. To make the fairest comparison with the strategic timing setting, we assume that the value of announcing correctly is 1 for player 1 and $\delta < 1$ for player 2. Thus being the second announcer is preferred if and only if $(13/16)\delta > 3/4$, or for fixed discount rate γ , if $r < (1/\gamma) \ln(13/12)$. The ex ante expected sum of utilities in the exogenous sequencing model is

$$U^{ex}(\delta) = 3/4 + (13/16)\delta.$$

The welfare properties under strategic timing also depend on δ , but in a more complicated way. Specifically, as δ decreases from 1 two things happen: (i) the probability that the players announce simultaneously, and hence no learning occurs, increases (from 0), and (ii) the game typically lasts longer, so greater delay costs are incurred. The exact welfare properties are difficult to work out in closed form for $\delta < 1$, so we consider the limiting case of welfare as $\delta \rightarrow 1$, since in this case we can easily derive explicit comparisons between the two models.

The probability that the first announcer (i.e., the one with higher signal quality) is correct can be worked out as $\Pr(Cor_1) = \Pr(Cor_1|A) = \Pr(p_1 \geq 1/2|A, |p_1 - 1/2| \geq |p_2 - 1/2|) = 5/6$. Since the second decision maker always makes the same announcement as the first, $\Pr(Cor_2) = 5/6$ as well. Thus we find that the difference in social welfare from endogenous timing is $(5/3 - (3/4 + 13/16))$ when there is no discounting, or almost 7%. The intuition is that, conditional on receiving a relatively uninformative signal, the exogenous sequencing forces a player to announce first with probability $1/2$, whereas when announcements are timed strategically, with high probability the poorly informed agent will learn from her opponent's signal and thus make a more informed announcement.

As δ decreases from 1 the expected welfare improvement from strategic timing increases, so that the difference of 7% as $\delta \rightarrow 1$ provides a lower bound on the social utility differential between the models. The reason is that as δ decreases in the exogenous sequencing model, the second decision maker is penalized with certainty, whereas when announcement times are chosen strategically, each player always chooses whether or not to delay, and is willing to delay

¹⁹ $\Pr(Cor_2) = \Pr(Cor_2|A) = \Pr(x_2 = a|A, x_1 = a)\Pr(x_1 = a|A) + \Pr(x_2 = a|A, x_1 = b)\Pr(x_1 = b|A) = \Pr(p_2 > 1/4|A)\Pr(p_1 > 1/2|A) + \Pr(p_2 > 3/4|A)\Pr(p_1 < 1/2|A) = (15/16)(3/4) + (7/16)(1/4) = 13/16$.

only if she expects to benefit from doing so. For example, if both agents are sufficiently well informed, then both announce at time zero, and no delay costs are incurred. Notice that when δ is sufficiently small (less than $2/3$), both agents will announce independently from each other at $t = 0$, regardless of their signals. Thus the ex ante probability that each is correct is $3/4$, and therefore $U^{en}(\delta) = 3/2$. In this case, $U^{ex}(\delta) \leq 3/4 + (13/16)(\frac{2}{3}) = 31/24 \approx 1.3$, so that the increase in welfare from endogenous timing is at least 15%. Finally observe that for $\delta < 2/3$ the expected social welfare in the endogenous timing model is constant, whereas it continues to decrease in the exogenous sequencing model. The welfare gains provided by the opportunity to strategically time announcements always fall short of full efficiency for $\delta < 1$ since either a delay cost is incurred by at least one agent, or else both agents announce independently at time 0, in which case no social learning takes place. Yet the limit point of the equilibrium outcomes as δ tends to one is the first best outcome.

A different basis of comparison between the models is to ask how frequently both agents announce the *pooled information state*, that is, the optimal choice if all private information were aggregated prior to announcing. Clearly, this happens with probability one under strategic timing in the limit $\delta \rightarrow 1$, since the agent with the more informative signal always announces “first” in equilibrium. On the other hand, in the worst case scenario for the endogenous timing model, δ is small enough that all types announce at 0, no learning occurs, and the pooled information state is announced by both agents with probability $(3/4)^2 + (1/4)^2 = 5/8$. In the exogenous sequencing model, in contrast, this happens about 79% of the time independently of δ .²⁰ Even though whenever δ is small enough that both agents announce the pooled information state more frequently in the exogenous sequencing model, the effect of greater delay costs dominates this informational advantage, so that expected welfare remains higher under strategic timing.

7 Extensions

7.1 Continuous Time

Most of the results extend to a continuous time setting in the following way. Consider instead of G^r a game played in continuous time with “reaction lags”. That is, players may make an

²⁰Note that the pooled information state is A if and only if $p_2 > 1 - p_1$. Thus one needs to compute $\Pr(p_1 > 1/2, p_2 > 1 - p_1) + \Pr(p_1 < 1/2, p_2 < 3/4 | p_2 < 1 - p_1)$ conditional on ω .

announcement at any time $t \geq 0$, but there is a strictly positive finite time that it takes to observe, process, and react to an announcement by another player. By a slight abuse of notation, call this parameter $r > 0$. Thus if j makes an announcement at some time t_j , i can not incorporate that knowledge into her own decision until time $t_j + r$. If i announces between t_j and $t_j + r$, it must be because she had already planned to do so, for she did not have time yet to react to j 's announcement. In such a case i makes her decision independently of j 's announcement. Denote by Γ^r the continuous time game so defined with reaction lag r .

The next lemma states that in any equilibrium of Γ^r , agents will announce only at the discrete times T^r that are allowed in the corresponding discrete time game G^r .

Lemma 1 *Given $\gamma > 0$, fix $r > 0$ and consider any PBNE σ^* of Γ^r . Then T^r has full measure under $\sigma_i^*(p)$ for all $p \in [0, 1]$, $i = 1, 2$.*

Proof. We claim that in any PBNE of Γ^r , agents will make decisions only at times kr , $k = 0, 1, \dots$ with probability one. Consider times $\tau \in (0, r)$. Any strategy that places positive probability on announcing at such times is strictly dominated by the strategy that is identical to σ_i^* except for shifting that mass to time $t = 0$, since conditional on being correct, the payoff at 0 is strictly greater, and, by hypothesis, no information can be acquired by any $\tau < r$. Therefore beliefs must not change before time r . Thus if an agent does not make an announcement at $t = 0$, she will wait at least until $t = r$. Next consider times $\tau' \in (r, 2r)$. Announcing at such times with positive probability is strictly dominated by announcing at r instead. Given that agents will not announce at times $\tau \in (0, r)$, no information will arrive in $(r, 2r)$, implying that beliefs are again constant on this interval, and conditional on being correct, the payoff at r is strictly greater than at any time after r . Thus if an agent has not announced at r , she will wait at least until $2r$. The claim follows in this fashion by induction on k . Thus in Γ^r agents will announce only at times that are allowed by G^r . ■

Lemma 1 establishes a connection between the games G^r and Γ^r which can be extended to show that the equilibrium outcomes in discrete and continuous time coincide for any $r > 0$.

7.2 Costly Information

Thus far we have presumed that agents are endowed with private information, and that the quality of the information varies randomly across agents. Further, we assumed that agents use

this information strategically in choosing when to make an announcement, and have observed that, in equilibrium, those with better information do better on average. This model then begs the question, “By what means do agents acquire their information?”. In this section, we suppose that information is available to the agents, but that there is a cost to gathering it, and that better information is more costly. One way to model this is to introduce an information source that sells signals to players. Suppose that before playing the “announcement game” G^r , players have the opportunity to simultaneously and privately purchase information. Any player who chooses not to buy information would begin the game with the initial prior $P(A)$. The information source offers a menu of contracts, selling signals of higher qualities at higher prices. Recall that $q = |p - 1/2| \in [0, 1/2]$ denotes the quality of a signal p . The menu takes the form of a function $c(q)$ with $c(0) = 0$, $c' > 0$ and $c'' \geq 0$. Conditional on purchasing a quality q , the signal indicates the correct state with probability $q + 1/2$. Thus signals of highest quality, $q = 1/2$, perfectly reveal the true state ω while signals of lowest quality, $q = 0$, which reveal no information, are equally likely to be correct or incorrect.

The value of signal quality is determined by the equilibrium welfare properties of the model. In particular, we must consider the expected utility of a player with signal quality q_i . Again consider the limiting case $\delta \rightarrow 1$ since it is the most tractable for this analysis, and the welfare properties of G^r approximate this limit for small r . In this case, the two players always make the same announcement, and this announcement is based solely on the better of the two pieces of information. Thus players’ welfare is determined by the quantity $\max\{q_1, q_2\}$. Signal quality can be viewed as a pure public good since both players benefit equally from either player purchasing a high quality signal. Given signal qualities, expected utilities are given by $u_i(q_1, q_2) = 1/2 + \max\{q_1, q_2\} - c(q_i)$.

If $c(q_i) = q_i$ it is easy to check that there are a continuum of asymmetric pure strategy Nash equilibria of the purchasing game, where one agent does not purchase information and the other agent purchases an arbitrary signal quality. There is also a unique symmetric Nash equilibrium, where both agents purchase no information. Note that the strategy of not purchasing ($q_i = 0$) is weakly dominant, and that it is a strict best response to any positive purchase by the other agent.

Allowing more general cost functions reduces the multiplicity problem, and alleviates the

perverse result that in the symmetric equilibrium, no information is purchased. In particular, assume $c'(0) < 1$, $c'' > 0$, and $c'(1/2) \geq 1$. Then there are exactly two asymmetric equilibria, where one agent chooses $q_i = 0$ and the other chooses $q_j = q^*$ given by $c'(q^*) = 1$. There is also a symmetric equilibrium where each agent mixes between $q_i = 0$ and $q_i = q^*$ with probabilities $c(q^*)/q^*$ and $1 - c(q^*)/q^*$, respectively. Thus the probability of purchasing information decreases with the cost of information, and increases with the optimal quality q^* . In all these equilibria information is under provided relative to the socially optimal arrangement where one agent purchases no information and the other purchases the quality characterized by $c'(q^o) = 2$.

7.3 Risk Aversion

We have assumed throughout the discussion that players are risk neutral. One interesting extension is to explore the implications of more general risk preferences. In the risk neutral case, the expected utility of announcing at time zero with belief π is $u(\pi) = \max\{\pi, 1 - \pi\}$. A risk averse agent has a utility function satisfying $\tilde{u}(\pi) < \max\{\pi, 1 - \pi\}$ for $\pi \in (0, 1)$, since the certainty equivalent to the lottery π is less than its expected value. We would like to be able to describe the qualitative effects that this has on the model's equilibrium predictions. Consider the case $n = 2$ for simplicity. The primary effect is that risk averse agents behave more patiently – they are willing to pay a greater cost in order to reduce uncertainty about the state variable. One can verify this directly by transforming the utility functions in the equation that defines the equilibrium cutpoints (see (9.2) in the Appendix), and resolving. Upon doing this, the result is that the cutpoints shift towards the boundaries, so that the maximum number of periods the game may last $\bar{k}(\delta)$ (weakly) increases. That is, the length of the game typically increases with risk aversion. In a particularly extreme form of risk aversion, $\tilde{u}(1/2) \rightarrow 0$. Yet even in this case, it can be verified that the game will end in finite time for any $\delta < 1$. Thus the main results do not depend on assuming a particular form of risk preferences. I have not verified that as δ tends to one that the game ends immediately for arbitrary utility functions, but I conjecture that this is true.

7.4 Other Utility Functions

The game we have analyzed has a specific payoff structure. One may wish to extend the model by allowing for more general utility functions. There are several interesting directions this could take. The first is to relax the maintained assumption of common interest. That is, one could assume that in addition to the common value component, there is an idiosyncratic private value of choosing either alternative. Indeed, most of economic theory presumes that choices reflect private preferences, rather than an attempt to identify a common objective. In a model of exogenous sequencing, Goeree, Palfrey, and Rogers (2003) explore such a setting, and find that the true state becomes known with certainty as the number of players diverges. A second possibility is to dispense with the pure information externalities environment. Note that doing so would complicate strategic aspects of the game, since it would introduce signalling issues. That is, when a player cares about how her announcement will effect the decisions of those who announce later, her optimal decision rule is much more difficult to characterize in general. One could assume that the utility of a player depends on the true state and the *group* decision (as defined by an appropriate voting rule), as opposed to the private decision.²¹ Such a model could, as a special case, maintain the common interest assumption, and capture situations such as committee decisions and jury verdicts. In this case, the motivation to announce quickly given a strong signal could arise either from the desire to signal this fact to those who follow, or as is the case above, from an explicit cost to waiting (or a combination of the two). Extending the model in this direction seems particularly interesting.

One could also imagine settings where individual interests are not perfectly aligned. This could arise in reputational models of information cascades,²² or in competitive environments, such as a patent race. The strategic timing aspects of these situations are more difficult to analyze, as there exists a tension between announcing correctly and quickly, and giving away information to others, which could be costly.

²¹See Hung and Plott (2001) for an experimental study of such an environment.

²²See, for instance, Scharfstein and Stein (1990) or Zwiebel (1992).

7.5 Relaxing Assumptions of the Model

The specification of the model as presented is somewhat restrictive. However, many of the results presented would still obtain under more general specifications. For instance, the assumption that the initial priors place equal weight on both states is not essential. Imagine instead that $P(A) \in (0, 1)$. Then the signal that results in a flat posterior is $p = 1 - P(A)$, rather than $p = 1/2$. The important part about a player's signal is her resulting posterior beliefs about the true state, which derive partially from her beliefs about the other players' information. Thus this generalization does not alter the fact that the game ends in finite time, and converges to zero length as the time interval vanishes. Note, though, that the definition of information-symmetry would have to be suitably modified. One natural way to do this is to phrase the definition in terms of posteriors rather than signals. In addition, the signal structure could also be suitably generalized. The essential feature of the conditional signal distributions, in addition to having convex support and well-behaved densities, is the property of first order stochastic dominance. That is, we require that higher types yield higher posterior beliefs for state A . It would be possible to prove analogues of the main results under such less restrictive assumptions. Note however, that any explicit comparison between endogenous timing and exogenous sequencing models, as in Section 6, would become much more difficult.

8 Conclusion

We have provided a framework in which it is convenient to analyze the properties of social learning where agents with differing signal qualities strategically choose the time at which to make their decisions. One nice feature of the model is a particularly simple and strong equilibrium characterization for the case of two players. The uniqueness of equilibria allows for unambiguous welfare comparisons to the standard model of exogenous sequencing. The primary finding is that expected social welfare is higher when timing is endogenous, and that this difference increases with the length of time between consecutive decisions. Another important finding is that, in sharp contrast to the standard model, information is used efficiently. There is no possibility of getting "stuck" in a situation where players may be choosing a suboptimal action despite the prevalence of useful information. Moreover this finding is not shared by similar models. For instance, in Zhang (1997), after the first agent has made an announcement, all agents decide

immediately with the same announcement, so that interesting dynamics of social learning may not be apparent.

Because of the necessity to allow instantaneous reactions in a continuous time model, the results of the discrete time model presented here may be more appropriate to institutional features that create lags. We may conclude that although a herd forms under strategic timing, the optimal timing decisions of agents are such that the herd is based on the *best* pieces of information in the economy, rather than on a random selection from the private information. This suggests that herds may be more often “right” than would be suspected from the model of exogenous sequencing. A final conclusion of the model may be drawn from the result that the information problem is solved asymptotically as n grows. Thus for large economies, one would expect that after an initial period where the best informed agents announce, at some point announcements occur very fast. This feature of the model may help to explain the dynamics of phenomena such as bank runs and currency crises.

9 Appendix

9.1 Beliefs Change only with Announcements

We claim that, given a signal p_i and information-symmetric strategies s , $\pi_i(p_i, t, h^t, s)$ is constant on all times t and histories h^t such that $x_j^\tau = 0$ for all $\tau < t$. Informally, an agent’s posterior beliefs about the true state can change only when the other agent makes an announcement. Note, however, that as time passes, each agent *does* update her beliefs about her opponent’s type, since given j ’s strategy s_j , the fact that the opponent has not yet announced restricts the support of $g(p_j|p_i)$.²³ In particular, let the updated distribution of j ’s type be $g(p_j|p_i, t, h^t, s)$. Thus

$$g(p_j|p_i, t, h^t, s) = \frac{g(p_j|p_i)}{\int_{\mathcal{S}(t, h^t, s)} g(p_j|p_i) dp_j} \quad (9.1)$$

where $\mathcal{S}(t, h^t, s)$ denotes the support of $g(p_j|p_i, t, h^t, s)$. By information-symmetry of s_j the set $\mathcal{S}(t, h^t, s)$ is symmetric about $1/2$. Thus it suffices to show that given a prior belief p_i , $\pi_i(p_i, t, h^t, s)$ is constant for all histories where neither agent has announced.

²³Recall that $g_i(\cdot|p_i, t, h^t, \sigma)$ denotes i ’s probability assessment of j ’s type given her own type, the time, the current history, and strategy profile σ , and that $g = (g_1, g_2)$.

Also let $\tilde{\pi}(p_1, p_2) = \frac{p_1 p_2}{p_1 p_2 + (1-p_1)(1-p_2)}$ denote the posterior on A given two conditionally independent signals p_1 and p_2 .

Using (9.1) we have

$$\begin{aligned}
\pi_i(p_i, t, h^t, s) &= \int_{\mathcal{S}(t, h^t, s)} \tilde{\pi}(p_i, p_j) g(p_j | p_i, t, h^t, s) dp_j \\
&= \left(\int_{\mathcal{S}(t, h^t, s)} g(p_j | p_i) dp_j \right)^{-1} \int_{\mathcal{S}(t, h^t, s)} \tilde{\pi}(p_i, p_j) g(p_j | p_i) dp_j \\
&= 2p_i \frac{\int_{\mathcal{S}(t, h^t, s)} p_j dp_j}{2p_i \int_{\mathcal{S}(t, h^t, s)} p_j dp_j + 2(1-p_i) \int_{\mathcal{S}(t, h^t, s)} (1-p_j) dp_j} \\
&= 2p_i \frac{\int_{\mathcal{S}(t, h^t, s)} p_j dp_j}{2 \left((2p_i - 1) \int_{\mathcal{S}(t, h^t, s)} p_j dp_j + (1-p_i) \int_{\mathcal{S}(t, h^t, s)} dp_j \right)} \\
&= p_i,
\end{aligned}$$

where the last equality follows since

$$\begin{aligned}
\int_{\mathcal{S}(t, h^t, s)} p_j dp_j &= (2p_i - 1) \int_{\mathcal{S}(t, h^t, s)} p_j dp_j + (1-p_i) \int_{\mathcal{S}(t, h^t, s)} dp_j \\
(1-p_i) \int_{\mathcal{S}(t, h^t, s)} 2p_j dp_j &= (1-p_i) \int_{\mathcal{S}(t, h^t, s)} dp_j,
\end{aligned}$$

which follows from the symmetry of $\mathcal{S}(t, h^t, s)$. Thus we have shown that when strategies are information-symmetric, agents learn nothing about the true state solely from the passage of time.

9.2 Proofs of Main Results

Proposition 1 *Fix $r > 0$ and $\gamma > 0$. G^r has an (essentially) unique PBNE s^* . Moreover, s^* is symmetric, information-symmetric, in pure strategies, and characterized by cutpoints.*

Proof. Define $\tilde{\pi}(p_1, p_2) = \frac{p_1 p_2}{p_1 p_2 + (1-p_1)(1-p_2)}$ to be the posterior belief that $\omega = A$ given two conditionally independent signals p_1 and p_2 . Also define $f(p_j | p) = 2(pp_j + (1-p)(1-p_j))$ to be i 's probability assessment of p_j given type p .

Given r and γ set $\delta = \exp(-\gamma r)$. Let $p_0 = 0$ and for $l > 0$ define p_l implicitly by

$$1 - p_l = \frac{\delta}{\int_{p_{l-1}}^{1-p_{l-1}} f(p|p_l) dp} \left(\int_{p_{l-1}}^{1-p_l} [1 - \tilde{\pi}(p_l, p)] f(p|p_l) dp + \int_{1-p_l}^{1-p_{l-1}} \tilde{\pi}(p_l, p) f(p|p_l) dp \right) \quad (9.2)$$

on the domain $(p_{l-1}, 1/2]$. It is straightforward to verify that the right hand side of (9.2) evaluated at $p_l = p_{l-1}$ is less than $1 - p_{l-1}$, and has derivative strictly greater than -1 on $(p_{l-1}, 1/2]$, implying that there is at most one solution. Let k denote the greatest l for which (9.2) has a solution. We claim that for any $r > 0$ the $\{p_l\}_{l=1}^k$ defined recursively by (9.2) are the cutpoints of the unique perfect Bayes-Nash equilibrium of G^r . Note that by the symmetry of the game, whenever a type $p < 1/2$ has a strictly dominated strategy to announce at some t , announcing at t is also strictly dominated for type $1 - p$. Thus we focus mainly on types $p < 1/2$, although it should be understood that the behavior of types greater than $1/2$ is also being described.

We first prove that for all $l = 1, \dots, k$, conditional on neither player having announced by time $(l-1)r$, all $p \in (p_{l-1}, p_l)$ (and all $p \in (1 - p_l, 1 - p_{l-1})$) have a strictly dominant strategy to announce immediately. To do so, we compare i 's expected payoff of announcing at $(l-1)r$ to the maximal continuation value of delay, i.e., the continuation value that obtains when j plays the continuation strategy that is optimal from i 's perspective. Specifically, this strategy is to announce at $(l-1)r$ if and only if $q_j > q_i$. Under this strategy, i learns all decision-relevant information immediately (at time $(l-1)r$), and is thus able to announce the pooled information state at the next time lr . Equation (9.2) compares the value of announcing at $(l-1)r$ to this maximal continuation value. As noted above, (9.2) has either zero or one solution on $p_l \in (p_{l-1}, 1/2]$. If it has no solution, then all types strictly prefer to announce at $(l-1)r$ since the payoff to announcing is greater than the maximal continuation payoff, and the game ends. If it has a solution p_l , then all types less than p_l have a strictly dominant strategy to announce at $(l-1)r$. Thus at $(l-1)r$, all types $p < p_{l-1}$ have already announced, and so they are eliminated from the support in (9.2).

We next prove that all types $\tilde{p} \in (p_l, 1/2]$ strictly prefer to delay at $(l-1)r$. To do so, we first show that for any such \tilde{p} , given that all $p \in (p_{l-1}, p_l)$ announce at $(l-1)r$, type \tilde{p} strictly prefers to delay at $(l-1)r$ provided all $p' \in (\tilde{p}, 1/2]$ delay at $(l-1)r$. We do this by comparing the value of announcing at $(l-1)r$ to a lower bound on the continuation value of delaying. Specifically, we use the value at the strategy “delay at $(l-1)r$, then announce at lr ” when j uses a “worst case” strategy. This lower bound obtains when j plays the strategy that *only* types $p < p_l$ announce at $(l-1)r$. In this case, we show that the expected payoff from delaying one period and then announcing at lr (which is a lower bound on the continuation value of delay at $(l-1)r$) in this

worst case scenario is greater than the value of announcing at $(l-1)r$. Consider the following inequality which states this formally.

$$1 - \tilde{p} < \frac{\delta}{\int_{p_{l-1}}^{1-p_{l-1}} f(p|\tilde{p})dp} \left(\int_{p_{l-1}}^{1-p_l} [1 - \pi(\tilde{p}, p)] f(p|\tilde{p})dp + \int_{1-p_l}^{1-p_{l-1}} \pi(\tilde{p}, p) f(p|\tilde{p})dp \right) \quad (9.3)$$

It is straightforward to verify that (9.3) is satisfied for all $\tilde{p} \in (p_l, 1/2]$.

This implies that for type $\tilde{p} = 1/2$ announcing at $(l-1)r$ is strictly dominated. Then since for every strategy s_j of agent j , $Eu_i(t_i, s_j|p_i)$ is continuous in p_i , there exists $\epsilon > 0$ such that announcing at $(l-1)r$ is strictly dominated by the strategy “wait at $(l-1)r$, announce at lr ” for all types $1/2 - \epsilon < p' < 1/2$. Let ϵ_1 be the supremum of all ϵ for which this holds. But then by (9.3) the fact that all types strictly above $1/2 - \epsilon_1$ delay implies $1/2 - \epsilon_1$ also delays. Then again by continuity, there exists a largest ϵ_2 such that all $1/2 - \epsilon_1 - \epsilon_2 < p'' < 1/2 - \epsilon_1$ delay. But then type $1/2 - \epsilon_1 - \epsilon_2$ has a strict preference to delay and we can take another open neighborhood to the left who strictly prefer to delay, etc. Define the monotonic sequence $p^n = 1/2 - \sum_{n=1}^n \epsilon_n$. Thus at each step, we know that all $p > p^n$ will wait. We claim that the p^n must converge to p_l . To show this, suppose instead that p^n converged to some $p^* > p_l$. Then for every $\epsilon > 0$ there exists an η such that $p^\eta - p^* < \epsilon$. Thus after a finite number η of rounds of iterative elimination of dominated strategies, all types $p \geq p^\eta$ wait at $(l-1)r$. Also note that by (9.3) if in addition all $p^* < p < p^\eta$ waited at $(l-1)r$ then p^* would have a strict preference to wait. But since $p^\eta - p^*$ is arbitrarily small, the probability that $p_j \in (p^*, p^\eta)$ is arbitrarily small, implying that for ϵ sufficiently small, player i with signal p^* strictly prefers to wait given only that types $p \geq p^\eta$ wait at $(l-1)r$.

To summarize, the $\{p_l\}$ defined by (9.2) define the cutpoints of the unique equilibrium of G^r . That is, at each date $(l-1)r$ all types $p < p_l$ have a strict preference to announce, and all types $p > p_l$ have a strict preference to delay. Type p_l is indifferent, and thus may choose to delay with an arbitrary probability. But since the set of indifferent types is of measure zero, their choices do not affect the expected utilities of other types. Thus the game has an essentially unique equilibrium, and is dominance solvable. ■

Proposition 2 *For every $\gamma > 0$ and $r > 0$, G^r ends in finite time. That is, there exists $k < \infty$ such that $s^*(p) < kr$ for all $p \in [0, 1]$ in the unique equilibrium s^* .*

Proof. Given the information-symmetry of s^* from Proposition 1, we confine attention to types $p \leq 1/2$. We can solve (9.2) explicitly for $p_l = p_l(p_{l-1}; \delta)$ to obtain the following.

$$p_l = 1/2 \left(1 - 1/\delta(1 - 2p_{l-1}) + \sqrt{(1/\delta - 1)(1 - 2p_{l-1})} \sqrt{1/\delta(1 - 2p_{l-1}) + 3 - 2p_{l-1}} \right). \quad (9.4)$$

Thus (9.4) provides a way to recursively solve for the cutpoint p_l , given only the previous cutpoint and the discount factor δ . Let $v(\delta) = 1/2(3 - 2/\delta)$. It is straightforward to verify that $p_l(p_{l-1}; \delta) > p_{l-1}$ on $[0, 1/2)$, is strictly increasing and concave on $[0, v(\delta)]$, and satisfies $p_l(v(\delta); \delta) = 1/2$ and $p_l(0; \delta) > 1/\delta - 1$, so that $\inf_{p_{l-1} \leq v(\delta)} p_l(p_{l-1}; \delta) - p_{l-1} = 1/\delta - 1$. Therefore, the number of cutpoints is bounded above by $\frac{\delta}{2(1-\delta)} < \infty$. ■

Before proving Theorem 1, we first prove a useful lemma.

Lemma 2 Let $\bar{k}(\delta) = \frac{\ln\left(\frac{(1/\delta - 1 + m(\delta))(3 - 2/\delta(1 + 1/\delta - m(\delta)))}{2(3 - 2/\delta)(-1/\delta + m(\delta))}\right)}{-\ln(1/\delta + m(\delta))}$, where $m(\delta) = \sqrt{(1/\delta - 1)(1/\delta + 3)}$. For every $\delta < 1$, $k(\delta) \leq \bar{k}(\delta)$.

Proof. Let $v(\delta) = \frac{3\delta - 2}{2\delta}$. Consider the line $h(p; \delta)$ through $(0, p_l(0; \delta))$ and $(v(\delta), 1/2)$. By concavity of $p_l(\cdot; \delta)$, $h(p; \delta) - p_l(p; \delta) < 0$ on $(0, v(\delta))$. The vertical distance from the 45° line to $h(p; \delta)$ is $d(p; \delta) = c(\delta)p + p_l(0; \delta)$, where $c(\delta) = \frac{1/\delta - m(\delta)}{3 - 2/\delta} - 1 \in (-1, 0)$. Consider the sequence $\{z_k\}$, where $z_{k+1} = z_k + d(z_k)$ and $z_0 = 0$. This sequence would describe the successive cutpoints if instead of using $p_l(p; \delta)$ we use the line $h(p; \delta)$ to compute cutpoints. Because $h(p; \delta)$ lies strictly below $p_l(p; \delta)$ on $(0, v(\delta))$, it follows that the first cutpoint z_k defined by $h(p; \delta)$ to lie above $1/2$ occurs no earlier than the first *true* cutpoint p_k , defined by $p_l(p; \delta)$ to do so. Using the linearity of $d(\cdot)$, one can solve the recurrence relation to obtain $z_k = \frac{p_l(0; \delta)}{c(\delta)} ((c(\delta) + 1)^k - 1)$. The claimed upper bound $\bar{k}(\delta)$ is derived by solving the equation $z_k = 1/2$ for k . ■

Theorem 1 For every $\gamma > 0$, G^r ends immediately as the time interval vanishes. That is, $\lim_{r \rightarrow 0} rk(\delta) = 0$.

Proof. Fix $\gamma > 0$. By Lemma 2, it is enough to show that $\lim_{\delta \rightarrow 1} \frac{-1}{\gamma} \ln(\delta) \bar{k}(\delta) = 0$. Applying L'Hopital's Rule and re-evaluating the limit proves the result. ■

References

- Banerjee, A., 1992. A Simple Model of Herd Behavior. *Quarterly Journal of Economics* 107, 797-817.
- Bikhchandani, S., Hirshleifer, D., Welch, Ivo., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades. *Journal of Political Economy* 100, 992-1026.
- Caplin, A., Leahy, J., 1994. Business as Usual, Market Crashes, and Wisdom After the Fact. *American Economic Review* 84,548-565.
- Chamley, C., 2004. Delays and Equilibria with Large and Small Information in Social Learning. *European Economic Review* 48(3), 477-501.
- Chamley, C., Gale, D., 1994. Information Revelation and Strategic Delay in a Model of Investment. *Econometrica* 62,1065-1085.
- Dasgupta, A., 2000. Social Learning with Payoff Complementarities. Working Paper: Yale University.
- Dasgupta, P., Maskin, E., 1986a. The Existence of Equilibrium in Discontinuous Games, I: Theory. *Review of Economic Studies* 53, 1-26.
- Dasgupta, P., Maskin, E., 1986b. The Existence of Equilibrium in Discontinuous Games, II: Applications. *Review of Economic Studies* 53, 27-42.
- Gale, D., 1996. What Have We Learned from Social Learning? *European Economic Review* 40, 617-628.
- Gul, F., Lundholm, R., 1995. Endogenous Timing and the Clustering of Agents' Decisions. *Journal of Political Economy* 103, 1039-1066.
- Goeree, J., Palfrey, T., Rogers, B., 2005. Social Learning with Private and Common Values. *Economic Theory*, forthcoming.
- Hung, A., Plott, C., 2001. Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions. *American Economic Review* 91,1508-1520.
- Jackson, M., Simon, K., Swinkels, J., Zame, W., 2002. Communication and Equilibrium in Discontinuous Games of Incomplete Information. *Econometrica* 70(5), 1711-1740.

- Kapur, S., 1994. Technological Diffusion with Social Learning. Working Paper: BirkBeck College.
- Milgrom, P., 1979. A Convergence Theorem fo Competitive Bidding with Differential Information. *Econometrica* 47(3), 679-688.
- Milgrom, P., Weber, R., 1987. Distributional Strategies for Games with Incomplete Information. *Mathematics of Operations Research* 10, 619-632.
- Reny, P., 1999. On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games. *Econometrica* 67(5), 1029-1056.
- Scharfstein, D., Stein, J., 1990. Herd Behavior and Investment. *American Economic Review* 80, 465-479.
- Simon, L., Stinchcombe, M., 1989. Extensive Form Games in Continuous Time: Pure Strategies. *Econometrica* 57(5), 1171-1214.
- Smith, L., Sorenson, P., 2000. Pathological Outcomes of Observational Learning. *Econometrica* 68, 371-398.
- Zhang, J., 1997. Strategic Delay and hte Onset of Investment Cascades. *Rand Journal of Economics* 28, 188-205.
- Zwiebel, J., 1992. Corporate Conservatism, Herd Behavior and Relative Compensation. *Journal of Political Economy* 103, 1-25.