Who goes first? Strategic Delay and Learning by Waiting*

Peter Wagner†

April 15, 2016

Abstract

This paper considers a timing game in which asymmetrically informed agents have the option to delay an investment strategically to learn about its uncertain return from the experience of others. I study the effects of information exchange through strategic delay on long-run beliefs and outcomes. Investment decisions are delayed when the information structure prohibits the occurrence of informational cascades. When there is only moderate inequality in the distribution of information, equilibrium beliefs converge in the long-run, and there is an insufficient aggregate investment relative to the efficient benchmark. When the distribution of information is more skewed, than the poorly informed drive out the well-informed, leading to a persistent wedge in posterior beliefs and excess investment.

*I am indebted to Martin Osborne, Colin Stewart, Ettore Damiano, Pauli Murto, Sven Rady as well as participants at the European Econometric Society Meeting for valuable discussions, helpful comments and suggestions. Financial support by the German Science Foundation (DFG) through SFB/TR 15 is gratefully acknowledged.

†Department of Economics, University of Bonn, Germany

Contents

1 Introduction

In 1929, the young German physician Werner Forssmann secretly conducted a risky self-experiment. He inserted a narrow tube into his arm and maneuvered it along a vein unto his heart. The procedure, known as cardiac catheterization, constituted a revolutionary breakthrough in cardiology and later earned him the Nobel prize in medicine. Forssmann's main contribution was the proof that cardiac catheterization was safe to perform on humans. The basic methods for the procedure had already been developed decades earlier and successfully tested on animals. It was widely believed, however, that inserting any object into the beating human heart was fatal, and thus there was a need for someone to put this hypothesis to the ultimate test.

The story of Werner Forssman is of someone who took action in an environment of "wait and see", in which everyone hoped for the independent initiative of a volunteer who resolves some of the risks relating to an uncharted course of action. There is a broad spectrum of areas in which these volunteer mechanisms play a crucial role. Palfrey and Rosenthal (1984) report the case of MCI, a telecommunications company, who fought for commercial access to AT&T's telephone networks in the 1960s, facing substantial legal fees and significant risk. The legal procedure ended with a favorable ruling by Federal Communications Commission (FCC) requiring AT&T to enable third parties to access their networks. In the end, the ruling benefited not only MCI, but also a host of other companies that were not previously involved in the case.

Empirically, it is a well-established fact that people learn from the behavior and experience of their peers. Peer learning effects have been found for example in the diffusion of innovations among health professionals (Becker, 1970), the enrollment in health insurance (Liu et al., 2014), the diffusion of home computers (Goolsbee et al., 2002), stock market entry (Kaustia and Knüpfer, 2012) and the introduction of the personal income tax (Aidt and Jensen, 2009). In environments in which no formal institution or informal arrangement exists that coordinates exploratory activities, how efficient is it to rely on the initiative of volunteers, and how well does such a decentralized mechanism aggregate dispersed information?

To study this problem, I consider a stopping game with asymmetric distribution of information and a pure informational externality. In this game, each agent has the option to make an investment. The investment generates an unknown return that depends on an uncertain state of the world. At the beginning of the game, agents privately receive information about the state and then decide independently how long to wait before taking action. The first agent who makes the investment realizes the state-dependent payoff and thereby reveals the state to the remaining agents. Uncertainty about the return of the investment and payoff observability generate a second-mover advantage that provides agents with an incentive to free-ride on others' initiative.

I characterize the Bayes-Nash equilibria allowing for an heterogeneous distribution of information. The equilibria can be broadly classified into two types. Equilibria may end immediately with some agent's immediate investment if the information structure is capable of generating an informational cascade. If this is not the case, then all robust equilibria exhibit delay. In an equilibrium with delay, agents wait for a period of time before making their investment. The delay is driven by the agents' expectation that someone else might invest first. The duration an agent is willing to wait provides a noisy signal to others about the value of the investment. The agents strategic considerations therefore influence beliefs which in turn affect investment decisions.

I study the effects of information exchange through strategic delay on long-run run beliefs and outcomes, and compare these to the efficient benchmark. Equilibria with delay can exhibit two structurally very different long-run outcomes. When information is fairly equally distributed, the natural equilibrium benchmark is one in which beliefs converge over time. All agents eventually become pessimistic about the state and investments stop. This equilibrium generates too little investment in aggregate relative to the efficient benchmark. In contrast, when the distribution of information is more skewed, there can be a persistent wedge in posterior beliefs between well and poorly informed agents.

I find that when optimism is high and information is distributed very unevenly, then the poorly informed tend to drive out the well-informed. Intuitively, the belief of well informed agents is more strongly correlated with the state of the world than that of the poorly informed. This means in particular that when the state is low, then poorly informed agents are less cautious in their investments. The well-informed, who tend to be more pessimistic in the low state, have thus a higher incentive to wait. As a result, the poorly informed learn increasingly less from the better informed and do themselves reveal more of their own information. In the limit, agents with more accurate information become entirely passive while their belief remains above that of the less informed.

In this equilibrium, agents stop with certainty regardless of the state of the world. The results in the literature typically suggest the opposite: when a public good is provided through voluntary contribution, then it is provided for at a socially insufficient level, because no agent takes into account the value of his own contribution to others. However, this insight is obtained almost exclusively through the analysis of symmetric equilibria of models featuring symmetric agents. The present paper deviates from this narrow focus on symmetric environments, characterizing the equilibrium outcomes in a more general model that allows agents to differ with respect to their endowment with information.

The paper is related to the literature on voluntary contributions to discrete public goods. These papers consider the strategic interaction between agents who face the binary decision of whether to contribute to a public good or not, and in which the public good is provided if the number of participants exceeds a given threshold. Such a model was first analyzed by Palfrey and Rosenthal (1984) who characterize its Nash equilibria. Consistent with standard logic, they find that in the unique symmetric equilibrium there is an insufficient provision of the public good. There are several extensions to their model allowing for the presence of informational asymmetry. Bliss and Nalebuff (1984) consider endogenous timing of voluntary contributions to a discrete good in a "war of attrition" framework. In their model, agents are privately informed about their own cost, and thus agents learn about others' participation only, but not about an underlying common state of the world.

There is also a natural connection to the literature on social learning, following the seminal articles of Bikhchandani et al. (1992) and Banerjee (1992). These papers consider models in which agents are ordered in a fixed sequence and learn from previous agents' actions about the common payoff to some risky action. They show that private information and sequential decision making can lead to informational cascades in which agents ignore their own information and herd on a socially undesirable action. Informational cascades arise in my model in symmetric equilibria, but never when there is a strong informational asymmetry.

Somewhat more closely related to this paper is Chamley and Gale (1994), who propose a variant of the game with endogenous timing of actions. In their model agents have an incentive to delay their action strategically when they expect to obtain additional information from other agents' decisions. A similar mechanism is at work in the present model as well, but the strategic setup is nevertheless quite different. In their model it is really the sequentiality of actions that is important – delays occur in their model only when agents are restricted to act at discrete times. In my model, delay arises naturally as a consequence of an informational spill-over that results from payoff observability.

Informational spill-overs from payoff observability have been studied in the strategic experimentation literature starting with Bolton and Harris (1999) and Keller et al. (2005). In these papers a group of agents dynamically choose between two actions (i.e., the arms of a bandit) one of which yields a risky and the other a safe payoff. Payoffs are observable giving rise to free-riding among agents and inefficient levels of experimentation with the risky action in equilibrium. Indeed, I view my model as a version of such a game, in which choosing the risky action is immediately fully revealing.

A number of papers study versions of games of strategic experimentation with asymmetrically informed agents. Those include non-competitive models in which agents are privately informed about their cost of delay (Décamps and Mariotti, 2004) or in which they privately observe their own payoffs (Rosenberg, Solan, and Vieille, 2007; Murto and Välimäki, 2011). Another array of papers considers model of competitive experimentation in which agents are privately informed about the realization of a common state variable (Malueg and Tsutsui, 1997; Moscarini and Squintani, 2010). To the best of my knowledge there is no paper that considers a model with a pure informational externality in which agents are asymmetrically informed about a common state variable.

The paper is structured as follows. The model, definitions and basic assumptions and the equilibrium concept are introduced in Section 2. Equilibrium and existence results are presented in Section 3. Section 4 presents the main result. Section 5 includes a discussion of efficiency and comparative statics. Section 5 concludes.

2 Model

There is a set of agents $N = \{1, \ldots, n\}$ who face the option to invest into a project with an uncertain return. The return for each project depends on the realization of an unknown state of the world $\theta \in \{H, L\}$, where $H > 0$ is arbitrary and L is normalized to -1. At the outset, all agents believe that $\theta = H$ with probability $p_0 \in (0, 1)$. Each agent decides if and when to stop.

The timing of the game is as follows. After observing their signals, the agents enter the *preemption phase* in which they decide sequentially whether to preempt the game and realize payoffs immediately. Preemption allows agents to move sequentially at time zero without delay, which is essential for equilibrium existence and for establishing an appropriate efficient benchmark.[1] When no agent preempts the game, they enter a *waiting phase* in which each agent chooses a stopping time representing the time at which an agent invests if no other agent has done so beforehand. Denote agent i's action by $t_i \in [0,\infty] \cup \{-i\}$ where $t_i = -i$ represents the event that agent i preempts the game and $t_i \geq 0$ is his stopping time conditional on reaching t_i in the waiting phase. When $t_i = \infty$, agent i waits indefinitely. The payoff for each agent i is

$$u_i(t_i, t_{-i}, \theta) = \begin{cases} e^{-r\max\{t_i,0\}}\theta & \text{if } t_i = \min_j t_j \\ e^{-r\max\{\min_j t_j,0\}}\max\{\theta,0\} & \text{if } t_i > \min_j t_j \end{cases}.$$

At the outset, it is commonly known that each agent $i \in N$ is endowed with a signal $s_i \in [0,1]$ that is drawn from a distribution $F_{i,\theta}(\cdot)$ which we assume is differentiable, has full support and a bounded density. A strategy for agent i is a function $\sigma_i : [0,1] \to [0,\infty] \cup \{-i\}$ with left limits. A strategy profile $(\sigma_i)_{i\in N}$ is a Bayes-Nash equilibrium if $\sigma_i(s_i) \in \arg\max_t \mathbb{E}[u_i(t, \sigma_{-i}(s_i), \theta)|s_i]$ for every $s_i \in [0,1]$. W.l.o.g. we limit attention to equilibria with $\Pr(\sigma_i(s_i) = 0) = 0$ for all $i \in N$.

For a given strategy profile $(\sigma_i)_{i\in N}$, let $\tau(s) = \min_{i\in N}\sigma_i(s_i)$ be first stopping time among all agents. Further, define $s_i^+ = \inf\{s_i|\sigma_i(s_i) = -i\}$ to be the lowest signal such that agent i preempts the game. Similarly, let $s_i^- = \inf\{s_i|\sigma_i(s_i) < \infty\}$ be the lowest signal such that agent i stops in finite time. We define $\inf\emptyset = 1$ for the case that one of these sets is empty. Finally,

[1]Without preemption, equilibria may fail to exist when some agent stops at $t = 0$ with positive probability. Then other agents may prefer to wait for that agent to move first, but since there is no first instance after $t = 0$, a best response may not exist.

define

$$A(t) = \{i \in N \big| \exists s_i \in [0,1] : \sigma_i(s_i) = t\}$$

to be set of agents that are "active" at time t, i.e., the set of agents for whom there exists a signal $s_i \in [0,1]$ such that agent i stops at t after observing s_i.

We assume that signal distributions satisfy the Monotone Likelihood Ratio Property (MLRP), that is, the likelihood ratio $F'_{i,H}(s_i)/F'_{i,L}(s_i)$ is increasing in s_i for each agent i. We shall make two further assumptions to render the strategic interaction interesting.

Definition 1 (Optimism). *Agent i is* weakly optimistic *if* $\mathbb{E}[\theta|s_i = 1] > 0$ *and* strongly optimistic*, if* $\mathbb{E}[\theta|s_i = 0] > 0$*.*

An agent is weakly optimistic if he assigns a positive expected value to θ after observing his best signal. A strongly optimistic agent assigns a positive expected value to θ after *any* signal. Weak optimism is a necessary condition for this agent's participation, since an agent for whom the expected value of stopping is negative at the outset would never act in any equilibrium.

Assumption 1 (Initial Optimism). *All agents are weakly optimistic*

Next, we assume that there is aggregate uncertainty about the state of the world. By aggregate uncertainty we mean that there is a signal for each agent so that this agent prefers not to act for some realization of another agents' signals.

Assumption 2 (Aggregate Uncertainty). $E[\theta|s_i{=}0, s_j{=}0] < 0$ *for any* $i \neq j$*.*

The assumption of aggregate uncertainty is important, because we are interested in studying issues relating the aggregation of dispersed information. In particular, the purpose of our model is to assess how well equilibria process information relative to the efficient benchmark. Aggregate uncertainty ensures that other's private information does not only influence when an agent invests, but also if he invests at all.

By Bayes' rule, agent i's belief that the state is H after observing signal s_i but before the beginning of the game, is

$$\Pr(H|s_i) = \frac{p_0 F'_{i,H}(s_i)}{p_0 F'_{i,H}(s_i) + (1-p_0) F'_{i,L}(s_i)}.$$

Denote by $p_i(\hat{s}_i, \hat{s}_{-i})$ agent i's belief that the state is H after observing signal s_i and conditional on the event that each agent j observed a signal no higher than s_j. By Bayes' rule, this belief is given by

$$p_i(s_i, s_{-i}) = \frac{\Pr(H|s_i)(s_i)\prod_{j\neq i} F_{j,H}(s_j)}{\Pr(H|s_i)(s_i)\prod_{j\neq i} F_{j,H}(s_j) + \Pr(L|s_i)(s_i)\prod_{j\neq i} F_{j,L}(s_j)}.$$

Define the *stopping value* of agent i at the signal profile s_i and $s_{-i} = (s_j)_{j\neq i}$ to be

$$\hat{u}_i(s_i, s_{-i}) = p_i(s_i, s_{-i})H - (1 - p_i(s_i, s_{-i})).$$

In some cases, agents may be endowed with particularly informative signals, that dominates others' information in the following sense.

Definition 2 (Dominant signal)**.** *Let s_i^* be the signal solving $E[\theta|s_i^*] = 0$. Agent i's signal is dominant, if $E[\theta|s_i \leq s_i^*, s_j = 1] \leq 0$ for all $j \neq i$.*

In other words, a dominant signal for agent i is a signal such that knowing that agent i's stopping value is negative discourages even the most optimistic competitor.

We denote by α the likelihood ratio of the posterior probability that the state is H, conditional on each agent i's signal being below s_i. It follows from Bayes' rule that

$$\alpha(s_1, \ldots, s_n) = \frac{p_0}{1-p_0}\prod_{i=1}^{n}\frac{F_{i,H}(s_i)}{F_{i,L}(s_i)}.$$

MLRP implies that $F_{i,H}/F_{i,L}$ is increasing for each i (Eeckhoudt and Gollier,

1995) and thus α is increasing in each of its arguments.

Further, we denote by $\lambda_{i,\theta}$ the *reverse hazard rate* of agent i's signal distribution in state θ given by

$$\lambda_{i,\theta}(s_i) = \frac{F'_{i,\theta}(s_i)}{F_{i,\theta}(s_i)}.$$

We shall impose the following technical assumption on the distribution of signals in the low state.

Assumption 3. *For every $i \in N$, we have*

$$(1) \qquad -\infty < \lim_{s_i \to 0} \frac{F''_{i,L}(s_i)/F'_{i,L}(s_i)}{F'_{i,L}(s_i)/F_{i,L}(s_i)} < 1.$$

This assumption is a mild regularity condition on the curvature of $F_{i,L}$. Essentially, it says that the curvature close to zero is neither too small nor too large relative to its slope. The condition is needed to ensure the function is sufficiently well-behaved around zero. Note that it is *not* a restriction on the informativeness of signals, because the restriction applies to the distribution of signals in the low state only, while informativeness is governed by the relative distribution of signals across states.

Denote by h_i the *reverse hazard rate ratio* (RHR) for agent i at $s_i \in [0, 1]$ is defined as the ratio of reverse hazard rates, and given by

$$h_i(s_i) = \frac{F'_{i,H}(s_i)/F_{i,H}(s_i)}{F'_{i,L}(s_i)/F_{i,L}(s_i)}.$$

It is well known that MLRP implies $\lambda_{i,H} > \lambda_{i,L}$ and thus $h_i > 1$. The hazard rate ratio h_i and the likelihood ratio of the public posterior α allows us to decompose the public posterior belief about the state into the common component and a private component.

$$\frac{p_i(s_i, s_{-i})}{1 - p_i(s_i, s_{-i})} = \alpha(s) h_i(s_i).$$

Here, α represents a measure of the information about the state that is commonly available to all agents. The factor h_i represents the information that agent i holds privately and it provides a measure of divergence of an agent's private belief from the public belief. Using decomposition, we write the stopping value for each agent i as follows. We have

$$\frac{\hat{u}_i(s_i, s_{-i})}{1 - p_i(s_i, s_{-i})} = \alpha(s) h_i(s_i) H - 1 \tag{2}$$

The left-hand side shows the stopping value relative to the probability of the low state. It measures the relative gain from delaying investment. The right-hand side shows that this value differs across agents only through differences in their respective RHR.

3 Socially optimal stopping

In this section we introduce a notion of efficiency that addresses the question of how agents should behave in order to maximize welfare. Our efficiency benchmark entails the restriction that agents cannot communicate their private information prior to deciding when to stop. We can interpret it as the solution to the "team problem" in which agents choose their strategies collaboratively, before observing their signals, so as to maximize the sum of their payoffs. Comparing equilibrium outcomes with this benchmark allows us to isolate inefficiencies in the use of information resulting from strategic effects and exclude those inefficiencies that are the result of the way information is processed in equilibrium. Our notion of efficiency is as follows.

Definition 3. *A strategy profile* $(\sigma_i)_{i=1}^n$ *is efficient if it maximizes*

$$\mathbb{E}\left[\sum_{i=1}^{n} u_i(\sigma_i(s_i), \sigma_{-i}(s_{-i}), \theta)\right].$$

An efficient allocation never entails any delay, because any outcome that is feasible through delayed stopping in the waiting phase can be achieved without delay in the preemption phase. To see this, fix any strategy profile σ and define $E_i = \{s_i|\sigma_i(s_i) < \infty\}$ to be the set of all signals for agent i for which i stops in finite time. Denote by $E = E_1 \times \ldots \times E_n$ the set of all signal profiles for which *some* agent stops in finite time. We call E the stopping region of σ. Now, consider an alternative strategy profile, in which agent i preempts the game if and only if $s_i \in E_i$ and waits indefinitely otherwise. This strategy profile generates the same stopping region as σ without delay, and thus increases the sum of payoff whenever stopping is indeed socially desirable. Finding the efficient strategy profile thus means determining the stopping region E that maximizes the expected welfare $\mathbb{E}[v(s)|s \in E]$, where

$$v(s) = \Pr(H|s)nH - \Pr(L|s).$$

Because preemption decisions have to be made autonomously by each agent, each agent should preempt if the expected sum of payoffs is positive conditional on his own signal *and* on the event that each other agent does not preempt the game.

The stopping region for an efficient strategy profile is characterized by thresholds, one threshold $\hat{s}_i$ for each agent i. This follows from the monotone likelihood ratio property: if it is socially optimal for an agent to preempt when his signal is s_i, then it must also be socially optimal to do so for any signal $s_i' > s_i$ as the higher signal implies a higher expected welfare.

Proposition 1. *If* $\hat{\sigma}$ *is an efficient strategy profile, there is a profile of signal thresholds* $\hat{s} = (\hat{s}_1, \ldots, \hat{s}_n) \in [0,1]^n$ *such that* $\hat{\sigma}_i(s_i) = -i$ *if* $s_i \geq \hat{s}_i$ *and*

$\hat{\sigma}_i(s_i) = \infty$ *otherwise. If* $\hat{s} > 0$*, then the threshold profiles satisfies* $\tilde{v}_i(\hat{s}) \leq 0$ *for all* i *and* $\tilde{v}_i(\hat{s}) = 0$ *if* $\hat{s}_i < 1$*, where*

$$\tilde{v}_i(\hat{s}) = \mathbb{E}\left[v(s) \,\middle|\, s_i = \hat{s}_i, s_{-i} < \hat{s}_{-i}\right].$$

Efficient strategy profiles can be viewed as equilibria of a modified game in which all agents pursue the common objective of maximizing social welfare. In this modified game, each agent i takes as given the strategies of others and then chooses the socially optimal response based on the information available to him: his own signal and the event that no other agent preempts. The best response for all agents is to preempt whenever the social value of doing so, based on their subjective posterior belief, is positive. In equilibrium it must therefore be the case that, conditional on no agent preempting the game, everyone expects the social value to be non-positive.

Figure 1 illustrates efficient stopping graphically for the case of two agents. Each agent $i = 1, 2$ preempts if his signal lies above the threshold $\hat{s}_i$, where the profile $(\hat{s}_1, \hat{s}_2)$ is given by the intersection of their zero-payoff curves. Naturally, the agents could do better if they were to pool their information before deciding whether to stop. In our benchmark agents fail to stop at signal profiles that would generate positive expected welfare if they were to pool information (Area I) and they do stop at signal profiles, at which it would be socially preferable not to (Area II).

Interestingly, in some cases it is efficient to ignore an agent's private information entirely. This is possible if information is distributed in such a way, that one agent's decision not to preempt overpowers any good news of others. Suppose, for example, there are two agents whose signals are drawn from distributions satisfying $F_{1,H}(s) = F_{1,L}(s)^{\beta}$ and $F_{2,H}(s) = F_{2,L}(s)^{\gamma}$ where $\beta > \gamma > 1$. These signal distributions satisfy MLRP and the reverse hazard rate ratios are constants given by $h_1(s_1) = \beta$ and $h_2(s_2) = \gamma$, respectively.

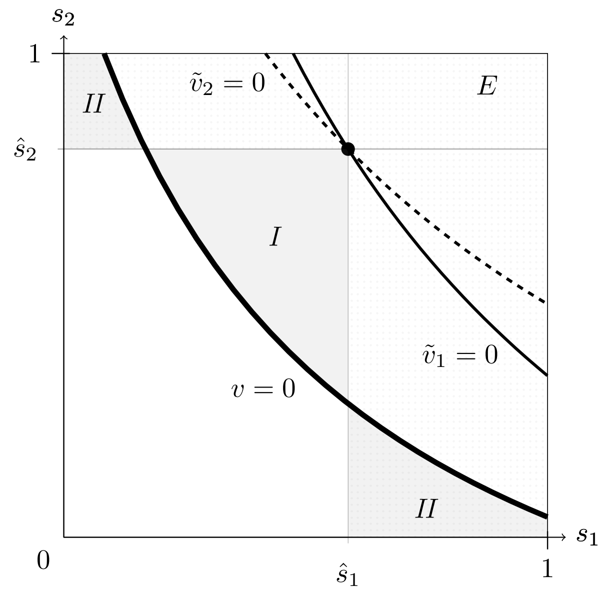

Figure 1

By the same logic as in Equation (2), we have the following inequality:

$$\tilde{v}_i(s_i) \leq 0 \Leftrightarrow \alpha(s_1, s_2)h_i(s_i) \leq 1/2H.$$

Since reverse hazard rate ratios are constant, the inequality cannot bind simultaneously for both agents. Therefore, by Proposition 1, the signal thresholds must be $\hat{s}_1 = 1/\sqrt{2\beta H}$ and $\hat{s}_2 = 1$. In this case, agent 2's information is entirely ignored, and agent 1's signal becomes decisive. What's happening intuitively is that agent 1's decision not to preempt is worse news than any potential good news that agent 2 may have. We can easily extend this logic to larger games by adding agents whose signal distributions are identical to that of agent 2. Taking this reasoning to the extreme yields a striking result: even as the number of agents becomes large and their information arbitrarily precise in aggregate, almost all of it can become irrelevant in the efficient benchmark under strong informational asymmetry.

4 Equilibrium Analysis

In this section we consider equilibrium outcomes of the model and discuss their properties. We begin with a preliminary result about the structure of equilibria which shows that equilibrium strategies are monotone and almost everywhere differentiable. We then provide a full equilibrium characterization for the case of two agents and generalize these to larger games with many agents. Finally, we consider the welfare implication and discuss connections to other literature.

4.1 Preliminaries

We begin by showing that equilibrium strategies are monotone and induce "smooth" distributions over stopping times. This result will be fundamental for the remaining analysis.

Proposition 2. *Let* $(\sigma_1, \ldots, \sigma_n)$ *be a Bayes-Nash equilibrium. Then for each* $i = 1, \ldots, n$*, we have:*

- (i) (Monotonicity) *Each* σ_i *is weakly decreasing with* $s_i^+ < 1$ *for at most one agent* $i \in N$*. If* $s_i^- < s_i^+$*, then* σ_i *is strictly decreasing on* (s_i^-, s_i^+)*.*
- (ii) (Smoothness) *Let* $s_i^- < s_i^+$*, and let* $D_i \subset (s_i^-, s_i^+)$ *be the set of discontinuities of* σ_i*. Then* σ_i *is differentiable on* $(s_i^-, s_i^+) \setminus D_i$*.*
- (iii) $|A(t)| \neq 1$ *on any open interval* $I \subset \mathbb{R}_+$*.*

Intuitively, the proposition says that each agent's equilibrium strategy is a decreasing function that has flat regions only at the upper and lower tail where it takes the values zero and infinity, respectively. If these flat regions do not meet, then there may be a countable number of downward jumps in the space between. Jumps in the equilibrium strategy of some agent i correspond to "passive" episodes in the equilibrium behavior of agent i, in the sense that there exists a time period during which agent i never stops for any of his signal realization. Discontinuities in the agents' strategies may arise as the result of changes in the set of actively participating agents.

Equilibrium strategies are monotone because agents who are more optimistic have a lower incentive to delay effort (this is the well-known cutoff property of (Fudenberg and Tirole, 1991). Intuitively, consider the trade-off of an agent choosing between stopping times t and $t' > t$. The gain from waiting at t until t' is equal to the expected loss avoided if another agent stops after t and before t' when the state is low. However, the agent incurs a loss from delay is decreasing in his signal. Thus, if an agent with signal s_i prefers to stop at t, the same holds for every signal $s_i' > s_i$.

Equilibrium strategies are "smooth" in the articulated sense because payoffs are differentiable with respect to stopping times and the signal distributions are well-behaved in the sense that they have full support with differentiable distribution functions. Therefore, small variation in signals leads a.s. to a small change in stopping times.

4.2 Two agents

In this section, we characterize the set of equilibria for the case of two agents. We differentiate between equilibria with *preemption* in which the game ends only in the preemption phase and equilibria with *delay* in which the game ends with positive probability in the waiting phase.

4.2.1 Equilibria with preemption

There are two reasons the game may end in the preemption phase. On reason is that an agent preempts the game because he has access to exceptionally accurate information and thus takes on the role of an informational leader whom others imitate. We call this scenario *informed preemption.* The second possibility is that a poorly informed, strongly optimistic agent preempts the game regardless of the realization of his signal, while all others wait for this agent wait for him to move. We refer to this second scenario as *uninformed preemption*.

Informed preemption. In an equilibrium with informed preemption, some agent i preempts the game if $\mathbb{E}[\theta|s_i] > 0$ and otherwise waits indefinitely. The other waits forever for sure. Informed preemption of agent i is possible in equilibrium if agent i's signal is dominant.

Proposition 3. *If agent i has a dominant signal, there exists an equilibrium with informed preemption by agent i.*

Informed preemption necessitates one agent to observe a dominant signal, so that the bad news that are conveyed through the agents' inaction at the beginning of the game "overpowers" any potential positive information the other agent might have. That the equilibrium conditions are satisfied follows immediately from the definition of dominant signals. The preempting agent expects that the other will never stop, and thus decides whether to preempt based only on his own information. If he preempts, the game is over. If he does not preempt, then the other agent updates his belief, and at this new belief, he assigns a negative expected value to the state by the definition of dominant signals. Thus it is optimal for him to wait indefinitely.

Uninformed preemption. The game may also end with certain preemption by a strongly optimistic agent. Certain preemption is optimal for an agent who is strongly optimistic, provided all other agents wait indefinitely, and waiting indefinitely is a best response for them to this one agent preempting for sure.

Proposition 4. *If agent i is strongly optimistic, then there exists an equilibrium with uninformed preemption by agent i.*

Uninformed preemption is conceptually more problematic than informed preemption. It is the only equilibrium in which the waiting phase is never reached, and thus our restriction to Bayes-Nash equilibria is less plausible. In particular, if we consider the analogous perfect Bayesian equilibrium of the fully dynamic equivalent of our game, then the existence of an equilibrium with uninformed preemption relies on the specification of off-equilibrium beliefs, and it is then not robust to slight perturbations to the payoff structure (Fudenberg et al., 1988). To see this point, suppose the preempting agent, agent 1 say, chooses to deviate and instead wait. How is the other suppose to respond? In equilibrium, agent 2 would have to wait indefinitely, even if he happens to be extremely optimistic himself. If we introduce a small change in payoffs, such that that there is a small probability that agent 2 prefers to never stop, waiting indefinitely is no longer a best response. The reason is that agent 2, after observing that agent 1 does not preempt, assumes that this is because agent 1 prefers to never stop. Thus agent 2's best response is to stop immediately thereafter. Naturally, given that agent 2 will respond this way, it is no longer optimal for agent 1 to preempt. The problem is that the equilibrium is sustained by action instead of information as is the case with informed preemption.

4.2.2 Equilibria with delay

In an equilibrium with delay the game ends with positive probability in finite time in the waiting phase. In such an equilibrium, each agent strategically delays taking action to take advantage of the possibility that another agent may move first. In this subsection, we show that the strategic interaction in these equilibria is captured by a pair of coupled differential equations. The long-run equilibrium outcomes correspond to fix-points of the associated dynamical system. Fix-points can exist in the interior of the space of signal profiles as well as on the boundary. We analyze equilibrium belief dynamics and illustrate how the location of fix-points and their stability attributes affect equilibrium properties.

When there are only two agents in the game, then it follows from Proposition 2, that if (σ_1, σ_2) is an equilibrium with delay, then σ_i is differentiable at $s_i < s_i^+$ for each $i = 1, 2$. Moreover, each agent's strategy has a differentiable monotone inverse, and thus we can use first-order necessary conditions to derive a system of differential equations whose solutions are candidates for inverse equilibrium strategies.

By monotonicity of the equilibrium strategies the distribution over agent i's stopping time in state θ can be written as $1 - F_{i,\theta}(\phi_i(t))$. Therefore, agent i's expected payoff from stopping at time $t > 0$ is given by

$$(3) \quad \Pr(H|s_i)\left(\int_0^t F'_{-i,H}(\phi_{-i}(\tau_{-i}))\phi'_{-i}(\tau_{-i})e^{-r\tau_{-i}}d\tau_{-i} + F_{-i,H}(\phi_{-i}(t))e^{-rt}\right)H$$
$$+ Pr(H|s_i)(1 - F_{-i,H}(\phi_{-i}(0))H - \Pr(L|s_i)F_{-i,L}(\phi_{-i}(t))\, e^{-rt}.$$

The first and second term is the expected payoff from taking action at t conditional on the state being high. Agent i with signal s_i assigns probability $\Pr(H|s_i)$ to this event. He receives payoff $e^{-r\tau_{-i}}H$ if agent $-i$ acts at $\tau_{-i} < t$, and otherwise he acts himself at time t and obtains the payoff $e^{-rt}H$. The third term represents the expected payoff if the state is low. In this case agent i receives a payoff of zero if the other agent acts before t, and otherwise he incurs a loss $-e^{-rt}$. Taking the first-order condition yields

$$r\Pr(H|s_i)F_{-i,H}(\phi_{-i}(t))H - r\Pr(L|s_i)F_{-i,L}(\phi_{-i}(t))$$
$$= -\Pr(L|s_i)F'_{-i,L}(\phi_{-i}(t))\phi'_{-i}(t).$$

Finally, substituting $s_i = \phi_i(t)$ and dividing both sides by the total probability of reaching time t, we can rewrite the last equation more succinctly as follows

$$(4) \qquad r\tilde{u}_i(\phi_i(t), \phi_{-i}(t)) = -\big(1 - p_i(\phi_i(t), \phi_{-i}(t))\big)\lambda_{-i,L}(\phi_{-i}(t))\phi'_{-i}(t).$$

Now, for any equilibrium $\sigma = (\sigma_1, \sigma_2)$ with delay, the pair of inverses (ϕ_1, ϕ_2)

must solve the system of differential equations

$$(5) \qquad \begin{aligned} -\phi_1'(t) &= Y_1(\phi_1(t), \phi_2(t)) \\ -\phi_2'(t) &= Y_2(\phi_1(t), \phi_2(t)) \end{aligned}$$

where

$$(6) \qquad Y_i(s_1, s_2) = \frac{r\tilde{u}_i(s_i, s_{-i})}{(1 - p_i(s_i, s_{-i}))\lambda_{i,L}(s_i)}.$$

By Proposition 2, strategies belonging to an equilibrium with delay must be monotonically decreasing, so that a solution path can belongs to an equilibrium if and only if it is strictly decreasing. Monotonicity and differentiability are in fact sufficient.

Proposition 5. *Let $s^+ = (s_1^+, s_2^+) \in [0,1]^2$ with $s_i^+ = 1$ for some i. Suppose ϕ is a pair of strictly decreasing inverse strategies solving* (5) *with initial condition $\phi(0) = s^+$. Then ϕ is an equilibrium.*

This result is a corollary to Proposition 8 which is proved in the appendix. To characterize the set of all Bayes Nash equilibria, we first find the fix points of the dynamical system (5) that are the solutions to the system of algebraic equations $Y_1(s_1, s_2) = Y_2(s_1, s_2) = 0$. The solutions lie along the zero-payoff curves which correspond to the set of all signal profiles at which an agent's stopping value is zero. Formally, the zero-payoff curve for agent i is defined as the set $\{(s_1, s_2) | \tilde{u}_i(s_i, s_{-i}) = 0\}$ of all signal profiles at which agent i's stopping value is zero. By the implicit function theorem, we can represent this set by a function φ_i, solving $\tilde{u}_i(s_i, \varphi_i(s_i)) = 0$ for each $i = 1, 2$. Note that, if $s_{-i} < \varphi_i(s_i)$, then $\tilde{u}_i(s_i, s_{-i}) < 0$ which implies $Y_{-i}(s_{-i}, s_i) < 0$, and thus $\phi_{-i}'(t) > 0$.

The path of a solution to the dynamical system is decreasing in the area above both zero-profit curves. Because each solution path will eventually converge to one of the fix points, a path belongs to an equilibrium only if it stays above these curves. We can interpret any point (s_1, s_2) in the diagram as a measure of the private information that remains with the agents. The closer s_i is to zero, the more information he has revealed to the other agent.

Interior limits. An interior limit is a fix point (s_1, s_2) of (5) with $s_i > 0$ for each i. It represents a long run equilibrium outcome in which each agent retains a positive amount of private information in the limit. At an interior limit, the reverse hazard rate is positive for each i, so that the denominator of each $Y_i(s_1, s_2)$ must be positive. Thus, by Equation (2), the point (s_1, s_2) must lie at intersection of the zero profit curves.

The following proposition shows that an interior limit exists if no agent has a dominant signal, and that any interior limit is also limit of an equilibrium with delay.

Proposition 6. *There exists an equilibrium that converges to an interior limit if both or neither agent has a dominant signal. Moreover, if* $s^* = (s_1^*, s_2^*)$ *is an interior limit and* $\varphi_i'(s_i^*)\varphi_{-i}'(s_{-i}^*) < 1$*, then there exists a unique equilibrium converging to* s^**.*

When no agent has a dominant signal, then, letting $\hat{s}_i = \varphi_i^{-i}(1)$, we have $0 = \tilde{u}_i(\hat{s}_i, 1) < \tilde{u}_{-i}(1, \hat{s}_i)$ for each i, and thus zero-payoff curves must indeed intersect. Similarly, when both agents have a dominant signal then $0 = \tilde{u}_i(\hat{s}_i, 1) > \tilde{u}_{-i}(1, \hat{s}_i)$. The stability properties of interior limits depend on the type of intersection. In general, when the zero-payoff curve for agent 1 intersects the zero-payoff curve for agent 2 from below (keeping s_1 on the horizontal axis) then the point of intersection is an unstable saddle point. Intuitively, at a point between these lines to the left of the intersection, the system flows upwards ($\tilde{u}_1 < 0$) and to the left ($\tilde{u}_2 > 0$), thus moving away from the point of intersection. In contrast, when the zero-payoff curve for agent 1 intersects the zero-payoff curve for agent 2 from above, then at a point between the lines to the left of the intersection, the system flows downwards ($\tilde{u}_1 > 0$) and to the right ($\tilde{u}_2 > 0$), thus moving towards the point of intersection.

The stability attributes of an interior limit determines the set of solution paths that converge to it. First, note that each Y_i is differentiable except potentially at the upper boundary when s_i approaches 1.[2] Therefore, the dynamical system (5) is locally Lipschitz in the interior and thus for any initial interior point s, there exists a unique solution. Now, starting at an interior limit s^*, we can choose any $s > s^*$ in a small neighborhood around the fix point, and solve (5) backwards in time starting at s. The solution is unique and strictly increasing, and by Rademacher's theorem we can extend the solution all the way to the boundary. The limit point then determines the initial signal pair (s_1^+, s_2^+). If s^* is an unstable saddle point, then there exists a unique solution path approaching s^* from above (i.e., the separatrix that runs from the boundary of the space of signal profiles along the crest to the saddle point).

Boundary limits. A boundary limit is a fix point (s_1, s_2) with $s_i = 0$ for one agent i. It represents a long run equilibrium outcome in which one agent stops with certainty in finite time, and by doing so perfectly reveals his private information. At a boundary limit, the reverse hazard for agent i goes to infinity, while for the other it must remain positive.[3] This implies that at boundary limit, the stopping value is zero for agent i, and positive for agent $-i$.

The following result shows that a boundary limit exists when an agent is strongly optimistic, and that for any boundary limit, there is a continuum of equilibria converging to it.

Proposition 7. *If agent i is strongly optimistic, then there exists a threshold*

[2] When $\lim_{s_i \to 1} F'_{i,L}(s_i) = 0$

[3] The latter follows from the fact that, by aggregate uncertainty (Assumption 2), the boundary limit cannot lie at the origin.

$\hat{s}_i$, such that for any $s_i^+ \leq \hat{s}_i$, there is an equilibrium that converges to a boundary limit.

The proposition tells us that when an agent is strongly optimistic, then there exists a continuum of equilibria converging to a boundary limit. The strongly optimistic agent must potentially preempt the game with positive probability, if that agent possesses "too much" information at the outset.

We construct such an equilibrium as follows. Suppose agent 1 is strongly optimistic. Let s_2 be the signal for agent 2 that solves $\varphi_2(s_2) = 0$. The signal s_2 has the property that the stopping value of agent 2 is zero if he observes s_2 and learns that agent 1 has received his worst signal. It is easy to check that $s^* = (0, s_2)$ is a boundary limit of the dynamical system (5).

The basic idea of the proof is to establish asymptotic stability of the boundary limit s^* and use this fact to show that there exists a continuum of strictly decreasing solution paths that converge to it. To this end, consider a sequence $(s_k)_{k\in\mathbb{N}}$ that converges to s^*, where $\alpha(s)h_1(s_k) = 1/H$ for all $k \in \mathbb{N}$ (a sequence moving to the boundary limit along the zero-payoff curve). For each k, solve the pair of coupled differential equations backwards in time with s_k as initial condition. Lipschitz continuity ensures a unique, strictly increasing solution path for each k that extends to the boundary of the space of signal profiles $[0, 1]^2$. Then taking the limit of these solution paths as $k \to \infty$ yields, in the limit, a strictly decreasing, continuous solution path that ends at s^*. If we now take a new point s along this path, and consider another point $s_\delta = (s_1 - \delta, s_2)$ with $\delta \in (0, s_1)$, then the solution path going through the newly selected point s_δ must also be strictly decreasing. For each δ, the point s_δ lies on a different solution path, and all of them (i) are strictly decreasing and (ii) converge to s^*.[4]

Figure 2 illustrates different types of equilibria for the case of two agents with symmetric signal distributions that have a monotone RHR. In each case,

[4]Note that when δ goes to s_1, the slope of the corresponding solution goes to infinity.

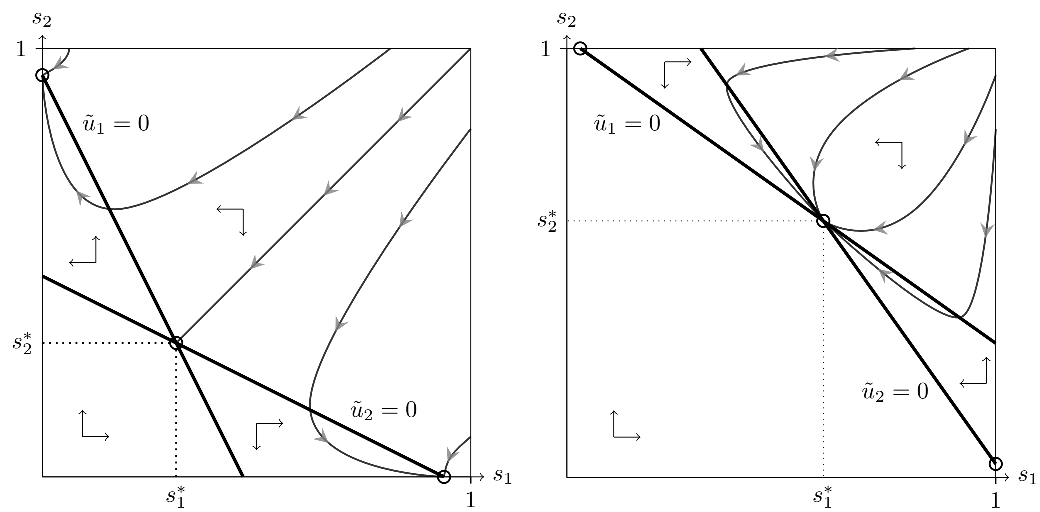

Figure 2: Phase diagrams for symmetric signal distributions with monotonically increasing RHR (left panel) and decreasing RHR (right panel).

there exists a unique interior limit. The left panel shows the phase diagram for the case in which the RHR is increasing and agents are strongly optimistic. In this case, there exist two equilibria with uninformed preemption, but neither agent has a dominant signal, and thus there is no equilibrium with informed preemption. Moreover, there exists a unique equilibrium with delay converging to an interior limit, and there is a continuum of equilibria converging to a boundary limit, one for each agent.

The right panel shows the phase diagram for the case in which the RHR is decreasing. When the RHR is decreasing, each agent has a dominant signal. Thus, there exist two equilibria with informed preemption, and multiple equilibria with delay converging to the unique interior limit. Decreasing RHR implies that neither agent is strongly optimistic, and thus there is no equilibrium with uninformed preemption, and no equilibrium converging to a boundary limit.

4.3 Many Agents

We now move on to consider games with more than two agents. The main insights from the case of two agents generalize to larger games. The essential properties of equilibria with preemption remain the same. Propositions 3 and 4 hold verbatim for any number of agents and the limitations for equilibria with uninformed preemption still apply. One difference is that there are stronger demands on a dominant signal, because the signal must informationally dominate all other agents' signal.

A substantial difference in larger games arises in the waiting phase, where agents can become now passive bystanders. With two agents, it is clear that delay is possible only if each agent stops with positive probability at every instant by Proposition 2. With more than two agents, any subset of at least two agents can engage in this sort of attrition game, allowing the others to wait and observe. This additional degree of freedom introduces an element of coordination into the game that substantially increases complexity.

To shed some light on source of this complexity, note at any instant, we can divide the set of all agents into those who are "active" in the sense that they stop with positive probability, and those who are "passive" in the sense that they stop with probability zero. Now, the inverse strategies for active agents at that instant are solutions to a system of differential equations obtained from the first-order conditions of active agents. The inverse strategies of passive agents are simply constants. The crucial observation is that the partition into active and passive agents is arbitrary and can in principle change at an arbitrary frequency as long as the probability that some active agent stops is such that it is indeed optimal for passive agents to wait. Because of the additional complexity, we do not attempt a full characterization of equilibria as in the two agent case. Instead, we focus on the characterization of equilibrium limit points.

Formally, periods of inactivity in the waiting phase correspond to jumps in an agent's stopping strategy. Because of these jumps, equilibrium strategies are generally not invertible. Instead, we work with the generalized inverse

$$\phi_i(t) = \sup\{s_i | \sigma_i(s_i) \geq t\}$$

which, for each i, gives the highest signal for which agent i stops after t. The function ϕ_i is the inverse of σ_i on its image, and its constant continuation elsewhere. Because σ_i is weakly decreasing and differentiable almost everywhere by Proposition 2, it follows that ϕ_i is weakly decreasing, continuous and almost everywhere differentiable. For convenience, we call the function ϕ_i an inverse strategy, and we say that a given profile $(\phi_1, \ldots, \phi_n)$ of inverse strategies constitutes an equilibrium if there exists an equilibrium $(\sigma_1, \ldots, \sigma_n)$ such that ϕ_i is the generalized inverse of σ_i for each $i \in N$.

By monotonicity of the equilibrium strategies, the distribution over agent i's stopping time in state θ can be written as $F_{i,\theta}(\phi_i(t))$. The probability that the earliest stopping time among all agents except i is after time t is equal to the joint probability that the signal of each agent $j \neq i$ is below $\phi_j(t)$, so that by conditional independence

$$G_{i,\theta}(t) = 1 - \prod_{j} F_{j,\theta}(\phi_j(t)). \tag{7}$$

Since ϕ_i continuous and almost everywhere differentiable, and each $F_{i,\theta}$ is differentiable and has full support, $G_{i,\theta}$ is continuous and almost everywhere differentiable.

$$\begin{aligned} \text{(8)} \quad \Pr(H|s_i) \left(\int_0^t e^{-r\tau_{-i}} dG_{i,H}(\tau_{-i}) + (1 - G_{i,H}(t)) e^{-rt} \right) H \\ + \Pr(H|s_i) G_{i,H}(0) H - \Pr(L|s_i)(1 - G_{i,L}(t))\, e^{-rt}. \end{aligned}$$

The interpretation is analogous to the two-agent case. The first and second term represent the expected payoff from taking action at t conditional on the state being high. Agent i with signal s_i assigns probability $\Pr(H|s_i)$ to this event. He receives payoff $e^{-r\tau_{-i}}H$ if another agent stops at $\tau_{-i} < t$, and otherwise he stops himself at time t and obtains the payoff $e^{-rt}H$. The second term represents the expected payoff if the state is low. In this case, agent i receives a payoff of zero if some agent stops before t, and otherwise he incurs a loss $-e^{-rt}$. The expectations are with respect to θ and τ_{-i}.

We follow essentially the same steps as in the two-agent case. A sufficient condition for agent i to be willing to delay stopping is that his marginal value of waiting is greater than zero:

$$(9) \quad -\Pr(H|\phi_i(t))(1-G_{i,H}(t))rH + \Pr(L|\phi_i(t))\left(G'_{i,L}(t)+r(1-G_{i,L}(t))\right) \geq 0.$$

Now, substituting he stopping distribution $G_{i,\theta}$ from equation (7) as well as

$$(10) \qquad G'_{i,\theta}(t) = -\prod_{j\neq i} F_{j,\theta}(\phi_j(t)) \left(\sum_{j\neq i} \frac{F'_{j,\theta}\left(\phi_j(t)\right)}{F_{j,\theta}\left(\phi_j(t)\right)} \phi'_j(t) \right).$$

Then, divide both sides of equation (9) by the total probability of reaching time t and substitute agent i's posterior belief p_i to obtain the following condition.

$$(11) \qquad r\tilde{u}_i(\phi_i(t), \phi_{-i}(t)) \leq -(1 - p_i(\phi_i(t), \phi_{-i}(t)) \sum_{j\neq i} \lambda_{j,L}(\phi_j(t))\phi'_j(t).$$

Consistent with intuition, the inequality tells us that an agent is willing to delay effort for an instant, only if the probability that some other agent will stop is higher than his the value he would receive if he were to stop.

The following result provides a sufficient condition for a profile of strategies to constitute a Nash equilibrium with delay.

Proposition 8. *A profile* $(\phi_1, \ldots, \phi_n)$ *of inverse strategies constitutes an equilibrium if the following hold.*

(i) *Every* ϕ_i *is continuous, differentiable a.e. and weakly decreasing.*

(ii) *For every* $i \in N$*, condition* (11) *holds at all* $t \geq 0$*.*

(iii) $|A(t)| \geq 2$ *for all* $t \geq 0$*.*

(iv) *For every* $i \in A(t)$ *and any* $t \geq 0$*, condition* (11) *holds with equality.*

(v) *If* $\phi_i(0) < 1$*, then* $\phi_j(0) = 1$ *for all* $j \neq i$*.*

The first property follows directly from the necessary conditions of Proposition 2. The second property says that there at least two agents active at any point in time in the waiting phase. The third property implies that at each time all agents at least weakly prefer to wait. The fourth property says that for any active agent, the strategy is pinned down by first-order conditions. The last property ensures that no more than one agent preempts the game. Note that there is no clear restriction on the choice of active agents, which introduces some degree of freedom.

For a given set of active agents, the inverse strategies are pinned down by property (*iv*). Using elementary operations and rearranging the equation system obtained by setting (11) equal for each $i \in A(t)$, we isolate the derivatives of inverse strategies of active agents. Doing so yields the dynamical system

$$(12) \qquad \begin{aligned} -\phi_1'(t) &= \mathbf{1}_{\{1\in A(t)\}} \cdot Y_1(\phi_1(t), \ldots, \phi_n(t)) \\ -\phi_2'(t) &= \mathbf{1}_{\{2\in A(t)\}} \cdot Y_2(\phi_1(t), \ldots, \phi_n(t)) \\ &\vdots \\ -\phi_n'(t) &= \mathbf{1}_{\{n\in A(t)\}} \cdot Y_n(\phi_1(t), \ldots, \phi_n(t)) \end{aligned}$$

where $\mathbf{1}_{\{i\in A(t)\}}$ is an indicator function that takes the value 1 if agent i is active and

$$(13)\ Y_i(s_1, \ldots, s_n) = \frac{r}{\lambda_{i,L}(s_i)} \left(\frac{1}{|A(t)|-1} \sum_{j\in A(t)} \frac{\tilde{u}_j(s_j, s_{-j})}{1-p_j(s_j, s_{-j})} - \frac{\tilde{u}_i(s_i, s_{-i})}{1-p_i(s_i, s_{-i})} \right).$$

In contrast to the two-agent case, not all agents may be active in the limit. Information asymmetry can result skewed posterior beliefs and agents with less accurate information, who tend to be more pessimistic than better informed ones, may eventually become too pessimistic to stop.

As a result, the specification of equilibrium limit points is more delicate than in the two-agent case. We must account for those agents who remain active in the limit. We thus define a equilibrium limit to be a profile $s^* = (s_1^*, \dots, s_n^*)$ such that there is a set $A \in N$, so that

$$\begin{aligned} Y_i(s_1^*, \dots, s_n^*) &= 0 \quad \forall i \in A \\ \tilde{u}_i(s_i^*, s_{-i}^*) &< \max_{j \in A} \tilde{u}_i(s_j^*, s_{-j}^*) \quad \forall i \notin A. \end{aligned} \tag{14}$$

In other words, an equilibrium limit is a fix-point of the dynamical system (12), restricted to a set $A \subseteq N$ of active agents, together with the requirement that the stopping value for every agent not in A is no higher than that of any active agent. The latter requirement makes sure than inactivity is in fact the result of inferior information. It is easy to see that, for given signal profile s and set A, condition (11) cannot hold along any solution path approaching s when the inequality is violated.

Analogously to the two-agent case, we call a limit point an interior limit if $s_i^* > 0$ for all i, and boundary limit otherwise. In the following, we provide a characterization of the stopping values at interior and boundary limits that provide some basic insights into the nature of long-run equilibrium outcomes.

Proposition 9. *Suppose $s^* = (s_1^*, \dots, s_n^*)$ is a limit point satisfying* (14) *for some $A \subset N$. Then the following holds.*

(i) If s^* is an interior limit, then $\tilde{u}_i(s_i^*, s_{-i}^*) = 0$ for all $i \in A$.

(ii) If s^* is a boundary limit, then there exists a unique $i \in A$ such that $s_i^* = 0$. Moreover, we have $\tilde{u}_i(s_i^*, s_{-i}^*) = 0$ and there is $u^* > 0$ such that $\tilde{u}_j(s_j^*, s_{-j}^*) = u^*$ for all $j \in A \setminus \{i\}$.

Interior limits generate a form of symmetry among agents in the sense that their stopping value at the limit is zero for each active agent. In particular, this implies that their posterior beliefs must be the same, that is, $p_i(s_i, s_{-i}) = p_j(s_j, s_{-j})$ for all $i, j \in A$. At a boundary limit, only the agent i with $s_i^* = 0$ has a stopping value of zero. The stopping value of all other agents equalizes as in the interior limit case, but may remain positive. This result is easy to see for the case of two agents, but the proposition generalizes to any higher number of agents.

To see that the stopping value of all agents must be zero at any interior limit s^*, notice that for (13) to be equal to zero, the expression in paranthesis must vanish. It is easy to see that this is possible for all $i \in A$ only if their stopping values are the same. We can thus simplify the expression and obtain that $Y_i(s^*) = 0$ only if $\tilde{u}_i(s_i^*, s_{-i}^*) = 0$ for all $i \in A$.

That a boundary limit can lie on the boundary for at most one agent follows immediately from the assumptioon of aggregate uncertainty. Recall that this assumption says that pooling the worst information of any two agents results in a negative stopping value for both of them.[5] The stopping value for the agent with $s_i^* = 0$ is zero because, in the limit, the stopping value and posterior beliefs of all active agents except i must equalize. Thus, for $j \in A \setminus \{i\}$, it follows again from $Y_j(s^*) = 0$ that $\tilde{u}_i(s_i^*, s_{-i}^*) = 0$. The remaining active agents retain a positive amount of private information and thus a positive stopping value.

We can use these facts to establish existence results that extend the statements of the two-agent case as follows.

Proposition 10. *The following holds.*

(*i*) *An interior limit exists if no agent has a dominant signal.*

[5] We made this assumption to ensure that agents are sufficiently interested in each others' information. Without this restriction, there would be equilibria that converge to the boundary in which the agents whose boundary is reached receice a positive stopping value.

(ii) A boundary limit s^* with $s_i^* = 0$ exists if agent i is strongly optimistic.

The existence results are analogous to the two-agent case. If no interior limit exist, then it must be the case that one agent is more optimistic than all other agents. Thus, this agent must have a dominant signal. Equilibria with preemption and delay are generally complementary, in the sense that if one type of equilibrium does not exist, then there must be an equilibrium of the other type. For the existence of a boundary limit strong optimism of some agent is enough. The remaining active agents either converge to a point that yields them with the same payoff, or, if such a point does not exist, all but one become eventually passive.

So far, we have only considered equilibrium limit points, but we have not actually contructed the associated equilibria. PRoviding a full characterization is difficult. The following example illustrates the various different equilibria that can arise.

4.4 Discussion

In this section we compare equilibrium outcomes to the efficient benchmark and detail how strategic incentives to delay investments affect long-run run equilibrium beliefs and outcomes. Inefficiencies arise because of a free-ride incentive. Agents choose to wait if they expect someone else to move first in an attempt to avoid the loss in a low state. We find that free-riding coupled with passive learning through delay can lead to a surprising reversal of the standard result saying that free-riding leads to insufficient investments: *when information is distributed unevenly, the less informed may invest too much relative the first best.*

Efficiency. In the efficient benchmark, the game always ends with preemption. Preemption occurs in equilibrium either through informed or uninformed preemption. Informed preemption is possible if some agent has access to a piece of critical information. In this case, the equilibrium structure coincides with the efficient benchmark in the sense that there is only one agent who stops with positive probability. However, this agent disregards the social value of the information he generates through his investment, and thus invests too little. In contrast to informed preemption, uninformed can generate excess investment. However, as we argued before, the equilibrium is problematic because it is not robust to slight perturbations of payoffs.

The plausible alternative to equilibria with informed preemption when no agent observes a dominant signal is an equilibrium involving strategic delay. When information agents have access to similarly accurate information, the natural equilibrium benchmark is one that converges to an interior limit. The aggregate stopping region of equilibria with delay that converge to an interior limit resemble the efficient stopping region. Both are determined by a sequence of threshold which is the intersection of zero-payoff lines, except that in equilibrium, agents disregard the social value of their actions. Consequently, there is insufficient investment in equilibrium.Moreover, note that the set of agents that are active in equilibrium might differ from those that would take part in the efficient benchmark which might drive an additional wedge between the social optimum and the long-run equilibrium outcome.

When the distribution of information is more skewed, and poorly informed agents actively participate along equilibrium path, their inferior information crowds out the activity of those agents who are better informed. Intuitively, the behavior of well-informed agents is more strongly correlated with the state of the world, than that of agents whose signal is noisy. This means in particular that when the state is low, then poorly informed agents are less cautious and due to their relatively higher rate of stopping the better informed become more passive. As a result, the poorly informed learn little from the better informed agents' delay and do themselves reveal more of their own information. In the limit, agents with more accurate information become entirely passive. while retaining a positive stopping value. In contrast, a less informed agent who remains active in the limit has a stopping value of zero in the limit (see Proposition 9).

There is a deeper insight that can be gained from our analysis: in the presence of indirect learning through delay, better information reduces the strategic incentive to stop. That this is the case can be seen immediately by closer inspection of (12). An active agent who has very accurate private information, as measured by the divergence of his private from the public belief, stop more slowly than an agent with who is endowed with less informative signals. More specifically, better information is associated with a larger value of his reverse hazard rate ratio h_i, and the larger this value the lower the stopping rate of the agent.

Fundamentally, what follows is that free-riding does not necessarily result in insufficient investment, as is the standard result in the literature (Foster and Rosenzweig, 2010), but can also lead to excess when information is distributed very unevenly. The majority of the literature that deals with issues of free-riding focuses on symmetric equilibria. Our analysis shows that allowing for asymmetric environments significantly expands the set of possibilities.

Additionally, notice that when there is excess investment in equilibrium, then the cost of this excess is borne mainly by the poorly informed. In fact, the likelihood that poorly informed stop is larger in a low state, because well informed agents are overall more likely to stop late or never. Thus, as far as production of information is concerned, there is no "exploitation of the great by the small" (Olsen, 1965) but rather an exploitation of the small by the great.

Replicator Dynamics. There is a strong connection between the equilibrium learning dynamics and replicator dynamics frequently used in ecology and evolutionary game theory. The classic "war of attrition" framework, on which our model is based, has a natural connection to replicator dynamics, because both have their roots in theoretical biology as dynamic models of competition. Here, however, replicator dynamics arise naturally as the result of strategic behavior. Recall that replicator dynamics are captured by a number of coupled first-order differential equations, the replicator equations, that characterize the changes in composition of a population over time as a function of its payoff or "fitness" in relation to the population average. Here, we consider evolution of the composition of private information distributed across the society of agents, but the dynamics are characterized by equations with similar properties. To see the similarity, note that (13) is also a function of agent i's stopping value in relation to an average.

The connection becomes most clear in the case of two agents. There, the first-order condition yield differential equations that generates dynamics identical to models of two competing species (a special case of replicator dynamics model) which becomes apparent in the phase diagram shown in Figure 2 (see Hofbauer and Sigmund, 1998). A crucial feature of the competing species models is that, under sufficiently strong competition between species, coexistence of both species is possible only at an instable fix point. Any small imbalance that favors one species leads to its complete dominance and the eventual extinction of the other. Here, we observe the same basic effect but applied to the revelation of information through strategic delay: in equilibrium, and agent who has relatively less private information must reveal more of it by stopping at a higher rate.

Herding. There is also a connection to the literature on social learning and herding. There are a number of papers that show how observational learning can lead to herding behavior. There, agents who observe others' behavior, and may ignore their own information in favor of the inferences based on their actions. In this way, observational learning leads to what Bikhchandani et al. (1992) call an *informational cascade*, wherein an agent's action is independent of his private information.

Here, we consider an environment with observational in which agents do not only observe others behavior but also the consequences to taking action. The game ends once an agent makes the investment and reveals the state of the world, so that herding in our model is, by construction, possible only on inaction. An informational cascade occurs in equilibrium only with informed preemption. After the preempting agent chooses to wait, every other agent then revises their belief about the state downwards, and chooses not to invest, independent of their private signal.

Informed preemption is the only type of equilibrium that exhibits a form of herding that is consistent with the notion that agents . For equilibria with delay, each active agents behavior is clearly dependent on their own information. What

5 Conclusion

The objective of this paper was to reveal some of the mechanisms that underly environments in which information is dispersed and privately informed agents learn from others through their strategic delay. I characterize long-run equilibrium outcomes for information structures that do not allow for informational cascades, and I show that these equilibria are typically inefficient. Investments are insufficient when agents are evenly well informed, but may also be excessive when information is distributed unevenly.

The basic setup of the model has been kept purposefully simple to retain tractability. It is however natural to consider extensions. For example, first-mover advantage or second-mover advantage appear plausible in many applications, such as R&D competition. Such a change would create a bias among agents for action or inaction, depending on whether we consider first or second-mover advantages, but qualitatively the basic insights in this paper remain the same.

Another possibility would be to study how private information affects free-riding in richer model in which experimentation occurs over time contemporaneously with learning from others' action. We may view the current model as a reduced form game in which the stopping payoffs represent the continuation value in an extended game in which a second round is played after agent stops.

A Proofs

Proof of Proposition 1. Because of the monotone likelihood ratio property, expected payoffs are non-decreasing in signals. Therefore, if it is optimal to stop for a given signal s_i of some agent i, then it must also be optimal and thus the stopping region is indeed characterized by a profile of thresholds $\hat{s}$. The optimal threshold profile solves

$$\max_{(\hat{s}_1,\ldots,\hat{s}_n)} p_0 \left(1 - \prod_{i=1}^{n} F_{i,H}(\hat{s}_i)\right) nH - (1-p_0)\left(1 - \prod_{i=1}^{n} F_{i,L}(\hat{s}_i)\right)$$

The associated Lagrangian is

$$\mathcal{L}(\hat{s}_1,\ldots,\hat{s}_n) = p_0 \left(1 - \prod_{i=1}^{n} F_{i,H}(\hat{s}_i)\right) nH - (1-p_0)\left(1 - \prod_{i=1}^{n} F_{i,L}(\hat{s}_i)\right) + \sum_{i\in N} \rho_i(\hat{s}_i - 0) + \sum_{i \in N} \mu_i(1-\hat{s}_i)$$

The efficient threshold profile $\hat{s}$ solves the necessary conditions

$$p_0 \prod_{j\neq i} F_{j,H}(\hat{s}_j) F'_{H,i}(\hat{s}_i) nH - (1-p_0)\prod_{j \neq i} F_{i,L}(\hat{s}_j) F'_{i,L}(\hat{s}_i) = \rho_i - \mu_i.$$

together with the Kuhn-Tucker conditions $\rho_i(s_i - 0) = 0$ and $\mu_i(1-s_i) = 0$ and $\rho_i, \mu_i \geq 0$ forall $i \in N$.

If $\hat{s}_i \in (0,1)$, then $\rho_i = \mu_i = 0$, so that the right-hand side is equal to zero, and $\hat{s}_i$ satisfies

$$\frac{p_0}{1-p_0} \prod_{j\neq i}^{n} \frac{F_{j,H}(\hat{s}_i)}{F_{j,L}(\hat{s}_i)} \frac{F'_{i,H}(\hat{s}_i)}{F'_{i,L}(\hat{s}_i)} = \frac{1}{nH}.$$

If there exists $i \in N$ with $\hat{s}_i = 1$, then $\rho_i = 0$ and $\mu_i > 0$. Thus,

$$\frac{p_0}{1-p_0} \prod_{j \neq i}^{n} \frac{F_{j,H}(\hat{s}_j)}{F_{j,L}(\hat{s}_j)} \frac{F'_{i,H}(\hat{s}_i)}{F'_{i,L}(\hat{s}_i)} = \frac{1}{nH} - \frac{\mu_i}{(1-p_0)\prod_{j \neq i}^{n} F_{j,L}(\hat{s}_j) F'_{i,L}(\hat{s}_i)} \leq \frac{1}{nH}.$$

Finally, if there exists an $i \in N$ with $\hat{s}_i = 0$, then $\rho_i > 0$ and $\mu_i = 0$. That is the case only if

$$\frac{p_0}{1-p_0} \prod_{j \neq i}^{n} \frac{F_{j,H}(\hat{s}_j)}{F_{j,L}(\hat{s}_j)} \frac{F'_{i,H}(0)}{F'_{i,L}(0)} = \frac{1}{nH} + \frac{\rho_i}{(1-p_0)\prod_{j \neq i}^{n} F_{j,L}(\hat{s}_j) F'_{i,L}(\hat{s}_i)} \geq \frac{1}{nH}.$$

for each i. □

Lemma 1. *The distribution over stopping times of each agent i induced by an equilibrium strategy σ_i has no atom except for at most one agent at time $t = 0$.*

Proof. (1) *There is at most one agent whose distribution over stopping times has an atom at $t = 0$.* Suppose there are two agents i, j who stop at time zero with positive probability. Then for each signal s_i with $\sigma_i(s_i) = 0$, agent i would do stricty better by stopping at time $\epsilon > 0$, for ϵ small.

(2) *There are no atoms at $t > 0$.* Suppose to the contrary that there is an atom at $t > 0$. Then by standard arguments, it cannot be optimal for any other to stop at a time $t - \epsilon$, for $\epsilon > 0$ small. But then $\sigma_i(s_i) = t$ cannot be a best response for any signal s_i of agent i, contradicting the hypothesis that there is an atom at t. □

Lemma 2. *Equilibrium strategies are non-increasing.*

Proof. We show that equilibrium payoffs are submodular. Let $q(s_i) = \Pr(H|s_i)$. The payoff of stopping at time t for agent i with signal s_i is

$$U_i^*(t, s_i) = q(s_i) \int_0^t e^{-rz} dG_{i,H}(z)\, H + e^{-rt} u_i^*(t, s_i)$$

where

$$u_i^*(t, s_i) = q(s_i)(1 - G_{i,H}(t))H - (1 - q(s_i))(1 - G_{i,L}(t)).$$

Let $\Delta U_i^*(t, t', s_i) = U_i^*(t, s_i) - U_i^*(t', s_i)$. Then, for $t' > t$ and $s_i' > s_i$, we have

$$\begin{aligned}
&\Delta U_i^*(t, t', s_i') - \Delta U_i^*(t, t', s_i) \\
&= q(s_i') \left(\int_t^{t'} e^{-rz} dG_{i,H}(z) \right) H + e^{-rt'} u_i^*(t', s_i') - e^{-rt} u_i^*(t, s_i') \\
&\qquad - \left(q(s_i) \left(\int_t^{t'} e^{-rz} dG_{i,H}(z) \right) H + e^{-rt'} u_i^*(t', s_i) - e^{-rt} u_i^*(t, s_i) \right) \\
&= (q(s_i') - q(s_i)) \left(\int_t^{t'} e^{-rz} dG_{i,H}(z) \right) H \\
&\qquad + e^{-rt'} \big(u_i^*(t', s_i') - u_i^*(t', s_i) \big) - e^{-rt} \big(u_i^*(t, s_i') - u_i^*(t, s_i) \big).
\end{aligned}$$

We can now use that

$$\int_t^{t'} e^{-rz} dG_{i,H}(z) \leq \int_t^{t'} e^{-rz} dG_{i,H}(z) + \int_t^{t'} r e^{-rz} G_{i,L}(z)\, dz = e^{-rt'} G_{i,H}(t') - e^{-rt} G_{i,H}(t)$$

and substitute

$$u_i^*(t, s_i') - u_i^*(t, s_i) = (q(s_i') - q(s_i)) \big[(1 - G_{i,H}(t))H + (1 - G_{i,L}(t)) \big]$$

to obtain the inequality

$$
\begin{aligned}
&\Delta U_i^*(t,t',s_i') - \Delta U_i^*(t,t',s_i) \\
&\leq (q(s_i') - q(s_i)) \left(e^{-rt'} G_{i,H}(t') - e^{-rt} G_{i,H}(t)\right) H \\
&\qquad + e^{-rt'}(q(s_i') - q(s_i))\big[(1 - G_{i,H}(t'))H + (1 - G_{i,L}(t'))\big] \\
&\qquad - e^{-rt}(q(s_i') - q(s_i))\big[(1 - G_{i,H}(t))H + (1 - G_{i,L}(t))\big] \\
&= (q(s_i') - q(s_i))(e^{-rt'}((1 - G_{i,L}(t')) + H) - e^{-rt}((1 - G_{i,L}(t)) + H)) \\
&\leq (q(s_i') - q(s_i))e^{-rt}(G_{i,L}(t) - G_{i,L}(t'))H \\
&< 0.
\end{aligned}
$$

Thus U_i^* is submodular, so that by Topkis' Monotonicity Theorem we have that

$$\sigma_i(s_i) = \arg\max_t U_i^*(t, s_i)$$

is non-increasing in s_i. □

Lemma 3. *Let $\sigma = (\sigma_1, \ldots, \sigma_n)$ be an equilibrium, and let $\phi = (\phi_1, \ldots, \phi_n)$ be its generalized inverse. Suppose ϕ_i is strictly decreasing for all $i \in A \subseteq N$ on an interval $I = (t, t')$ with $t' > t > 0$. Then ϕ_i is differentiable on I for each $i \in A$.*

Proof. We show that $F_{i,L}(\phi_i(t))$ is Lipschitz-continuous for each i. Because $F_{i,H}$ has full support by hypothesis, it follows then that σ_i is differentiable almost everywhere.

By definition of U^* it follows that

$$\Delta U_i^*(t,t',s_i) = q(s_i)\int_0^t e^{-rz} dG_{i,H}(z)\, H + e^{-rt}u_i^*(t,s_i)$$
$$- q(s_i)\int_0^{t'} e^{-rz} dG_{i,H}(z)\, H - e^{-rt'}u_i^*(t',s_i)$$
$$= e^{-rt}u_i^*(t,s_i) - e^{-rt'}u_i^*(t',s_i) \; - q(s_i)\int_t^{t'} e^{-rz} dG_{i,H}(z)\, H.$$

Agent i prefers $t = \sigma_i(s_i)$ over $t' \in (t, \sigma_i(s))$, and therefore it must be the case that $\Delta U_i^*(t,t',s_i) \geq 0$. Thus, it follows

$$e^{-rt}u_i^*(t,s_i) - e^{-rt'}u_i^*(t',s_i) \; \geq q(s_i)\int_t^{t'} e^{-rz} dG_{i,H}(z)\, H. \tag{15}$$

We further have

$$\int_t^{t'} e^{-rz} dG_{i,H}(z) \geq e^{-rt'}\int_t^{t'} dG_{i,H}(z) = e^{-rt'}(G_{i,H}(t') - G_{i,H}(t)). \tag{16}$$

Using a zero-addition, we find that

$$e^{-rt}u_i^*(t,s_i) - e^{-rt'}u_i^*(t',s_i)$$
$$= e^{-rt}u_i^*(t,s_i) - e^{-rt'}u_i^*(t,s_i) + e^{-rt'}u_i^*(t,s_i) - e^{-rt'}u_i^*(t',s_i)$$
$$= (e^{-rt} - e^{-rt'})u_i^*(t,s_i) + e^{-rt'}q(s_i)(G_{i,H}(t') - G_{i,H}(t))H$$
$$- e^{-rt'}(1-q(s_i))(G_{i,L}(t') - G_{i,L}(t)).$$

where we used the definition of u_i^* in the last equation. Rearranging the last equality yields

$$(e^{-rt} - e^{-rt'})u_i^*(t,s_i) \tag{17}$$
$$= e^{-rt}u_i^*(t,s_i) - e^{-rt'}u_i^*(t',s_i) - e^{-rt'}q(s_i)(G_{i,H}(t') - G_{i,H}(t))H$$
$$+ e^{-rt'}(1-q(s_i))(G_{i,L}(t') - G_{i,L}(t)).$$

Now, use (15) and (16) successively to obtain

$$(e^{-rt} - e^{-rt'})u_i^*(t, s_i) \geq e^{-rt'}(1 - q(s_i))(G_{i,L}(t') - G_{i,L}(t)).$$

The exponential function e^{-rt} is Lipschitz-continuous on the positive real-line with Lipschitz bound r, and therefore $r(t' - t) \geq e^{-rt} - e^{-rt'}$. Althogher, it follows

$$L(t, t')(t' - t) \geq (G_{i,L}(t') - G_{i,L}(t))$$

where

$$L(t, t') = \frac{r}{e^{-rt'}} \frac{u_i^*(t, \phi_i(t))}{1 - q(\phi_i(t))}.$$

The function $L(t, t')$ is positive because ϕ_i is strictly decreasing on (t, t') and thus $u^*(t, \phi_i(t)) > 0$. Second, $L(t, t')$ is finite because $q(\phi_i(t')) < 1$ (if $q(\phi_i(t')) = 1$ agent i with signal $\phi_i(t')$ would not want to wait until $t' > 0$). Therefore, $L(t, t')$ is continuous and bounded on $I \times I$ which implies $L^* = \max_{(t,t') \in I \times I} L(t, t')$ exists. Hence,

$$|G_{i,L}(t') - G_{i,L}(t)| \leq L^* |t' - t|$$

for all $t \geq t'$ in I, which means that $G_{i,L}$ is locally Lipschitz-continuous. Moreover, for any $j \in A \setminus \{i\}$ we have

$$\begin{aligned} |G_{i,L}(t') - G_{i,L}(t)| &= \left| \prod_{l \neq i} F_{l,L}(\phi_j(t)) - \prod_{l \neq i} F_{l,L}(\phi_j(t')) \right| \\ &\geq \prod_{l \neq i,j} F_{l,L}(\phi_j(t')) \left| F_{j,L}(\phi_j(t)) - F_{j,L}(\phi_j(t')) \right|. \end{aligned}$$

In equilibrium, we have $\prod_{l \neq i,j} F_{l,L}(\phi_j(t')) > 0$, and thus we can combine the last two inequalities to obtain

$$|F_{j,L}(\phi_j(t')) - F_{j,L}(\phi_j(t))| \leq \frac{L^*}{\prod_{l \neq i,j} F_{l,L}(\phi_j(t'))} |t' - t|.$$

which implies that each $F_{j,L}(\phi_j(\cdot))$ is locally Lipschitz-continuous, as well. Now, each $F_{j,L}$ is strictly increasing and continuously differentiable by assumption, and hence it is invertible, and the derivative of the inverse $F_{j,L}^{-1}$ is again differentiable with bounded derivative (since $F_{j,L}$ has full support). Thus, $F_{j,L}^{-1}$ is Lipschitz-continuous with some Lipschitz-bound M, and

$$\begin{aligned}|\phi_j(t) - \phi_j(t')| &= |F_{j,L}^{-1}(F_{j,L}(\phi_j(t)) - F_{j,L}^{-1}(F_{j,L}(\phi_j(t'))| \\ &\leq M|F_{j,L}(\phi_j(t)) - F_{j,L}(\phi_j(t'))| \\ &\leq \left(\frac{ML^*}{\prod_{l\neq i,j} F_{l,L}(\phi_j(t))}\right)|t-t'|.\end{aligned}$$

The last inequality shows that ϕ_j is locally Lipschitz-continuous. Since this holds for all i, it follows from Rademacher's Theorem that every ϕ_j is differentiable almost everywhere on $\mathbb{R}_+$. □

Proof of Proposition 2. (i) By Lemma 2, equilibrium strategies are non-decreasing which implies $0 \leq s_i^- \leq s_i^+ \leq 1$. By Lemma 1, the distribution over stopping times of every agent has no atoms except at time zero. Therefore, $s_i^+ = 1$ for all agents except at most one.

(ii) Follows from Lemma 3.

(iii) Suppose $A(t) = \{i\}$ on some open interval $(t_0, t_1) \subset \mathbb{R}_+$. But agent i receives a strictly higher payoff from stopping at t_0 than at t_1 which implies that stopping at t_1 cannot be a best response. □

Lemma 4. *Let $n = 2$. There exists a fix-point in the interior if no agent has a dominant signal.*

Proof. It is sufficient to show that there are points (s_1, s_2) and (s'_1, s'_2) such that $u_1(s_1, s_2) < u_2(s_2, s_1)$ and $u_1(s'_1, s'_2) > u_2(s'_2, s'_1)$.

(i) Suppose neither agent is strongly optimistic. Let $\hat{s}_i$ be defined such that $\hat{u}_i(s_i, 1) = 0$. Since neither agent has a dominant signal, we have $\hat{u}_1(\hat{s}_1, 1) = 0 < \hat{u}_2(1, \hat{s}_1)$ and $\hat{u}_2(\hat{s}_2, 1) = 0 < \hat{u}_1(1, \hat{s}_2)$.

(ii) Suppose agent i but not agent $-i$ is strongly optimistic. Define $s_i^\dagger$ such that $\hat{u}_i(0, s_i^\dagger) = 0$. Then,

$$0 = \hat{u}_i(0, s_i^\dagger) = \mathbb{E}[\theta | s_i = 0, s_{-i} < s_i^\dagger] < \mathbb{E}[\theta | s_i = 0, s_{-i} = s_i^\dagger] = \hat{u}_{-i}(s_i^\dagger, 0).$$

(iii) If both agents are strongly optimistic, then we can apply (ii) for each of them. □

Proof of Proposition 3. We prove the result for $n \geq 2$. Suppose agent i has a dominant signal. Let s_i^* solve $\mathbb{E}[\theta | s_i] = 0$. Set

$$\sigma_i(s_i) = \begin{cases} -i & \text{if } s_i > s_i^* \\ \infty & \text{if } s_i < s_i^* \end{cases}$$

and let $\sigma_j(s_j) = \infty$ for all $j \neq i$. The payoff for agent i i

$$U_i(s_i) = \begin{cases} \mathbb{E}[\theta | s_i] & \text{if } s_i > s_i^* \\ 0 & \text{if } s_i < s_i^* \end{cases}$$

If $s_i < s_i^*$ agent i cannot gain by stopping at a finite time. If $s_i > s_i^*$ and agent i deviates by stopping at $t > 0$, then his payoff is $e^{-rt}\mathbb{E}[\theta | s_i] < U_i(s_i)$. No agent $j \neq i$ can gain by preempting before agent i. If agent $j \neq i$ chooses a stopping time $t \geq 0$, his payoff is

$$e^{-rt}\big(\Pr(H|s_j) F_{i,H}(s_i^*) H - \Pr(L|s_j) F_{i,L}(s_i^*)\big) < e^{-rt}\mathbb{E}[\theta | s_j, s_i < s_i^*] < 0.$$

Hence, this deviation is not profitable. □

Proof of Proposition 4. We prove the result for $n \geq 2$. Suppose agent i is strongly optimistic. Set $\sigma_i(s_i) = -i$ and $\sigma_j(s_j) = \infty$ for all $j \neq i$. The payoff for agent i is $U_i(s_i) = \mathbb{E}[\theta|s_i]$. By strong optimism, $U_i(s_i) \geq 0$ for all s_i. If agent i deviates by stopping to $t > 0$, his payoff is $e^{-rt}\mathbb{E}[\theta|s_i] < U_i(s_i)$. For any agent $j \neq i$, the payoff is $U_j(s_j) = E[\max\{\theta, 0\}|s_j]$ which is the maximum attainable payoff, so no deviation can be profitable. $\square$

Proof of Proposition 6. Define

$$e_i(x, y) = \alpha(x, y)h_i(y)H - 1$$

Let $s^* = (x, y)$ be an interior limit. The Jacobian for the dynamical system is

$$J = \begin{pmatrix} \frac{\lambda'_{L,1}(x)e_2(x,y)-\lambda_{L,1}(x)\partial_x e_2(x,y)}{\lambda_{L,1}(x)^2} & -\frac{\partial_y e_2(x,y)}{\lambda_{L,1}(x)} \\ -\frac{\partial_x e_1(x,y)}{\lambda_{L,2}(y)} & \frac{\lambda'_{L,2}(y)e_1(x,y)-\lambda_{L,2}(y)\partial_y e_1(x,y)}{\lambda_{L,2}(y)^2} \end{pmatrix}.$$

Note that if s^* is an interior limit, then $e_2(s^*) = e_1(s^*) = 0$. Thus, the Jacobian becomes

$$J = \begin{pmatrix} -\frac{\partial_x e_2(x,y)}{\lambda_{L,1}(x)} & -\frac{\partial_y e_2(x,y)}{\lambda_{L,1}(x)} \\ -\frac{\partial_x e_1(x,y)}{\lambda_{L,2}(y)} & -\frac{\partial_y e_1(x,y)}{\lambda_{L,2}(y)} \end{pmatrix}.$$

The associated characteristic polynomial is given by

$$\det(J - \rho I) = \left(-\frac{\partial_x e_2(x,y)}{\lambda_{L,1}(x)} - \rho\right)\left(-\frac{\partial_y e_1(x,y)}{\lambda_{L,2}(x)} - \rho\right) - \frac{\partial_x e_1(x,y)}{\lambda_{L,1}(x)}\frac{\partial_y e_2(x,y)}{\lambda_{L,2}(x)}.$$

The roots of the characteristic polynomial are

$$\rho_{1,2} = -\frac{\lambda_{L,1}(x)\partial_y e_1(x,y) + \lambda_{L,2}(y)\partial_x e_2(x,y)}{2\lambda_{L,1}(x)\lambda_{L,2}(y)} \pm \left(\frac{(\lambda_{L,2}(y)\partial_x e_2(x,y) + \lambda_{L,1}(x)\partial_y e_1(x,y))^2}{4\lambda_{L,1}(x)^2\lambda_{L,2}(y)^2} - \frac{4\lambda_{L,1}(x)\lambda_{L,2}(y)(\partial_x e_2(x,y)\partial_y e_1(x,y) - \partial_y e_2(x,y)\partial_x e_1(x,y))}{4\lambda_{L,1}(x)^2\lambda_{L,2}(y)^2} \right)^{1/2}.$$

By the implicit function theorem, the nullclines φ_1, φ_2 defined implicitly through $e_1(s_1, \varphi_1(s_1)) = 0$ and $e_2(\varphi_2(s_2), s_2) = 0$, have the slopes

$$\varphi_1'(x) = -\frac{\partial_x e_1(x,y)}{\partial_y e_1(x,y)}, \qquad \varphi_2'(y) = -\frac{\partial_y e_2(x,y)}{\partial_x e_2(x,y)}$$

If $\varphi_1'(s_1)\varphi_2'(s_2) < 1$, then

$$\partial_x e_2(x,y)\partial_y e_1(x,y) - \partial_y e_2(x,y)\partial_x e_1(x,y) < 0.$$

Thus, the Eigenvalues r_1 and r_2 have opposite signs, which implies that the interior steady is a saddle point, and hence unstable. Thus, there exists a unique trajectory (the separatrix) that converges to s^*, and this trajectory constitutes an equilibrium path. □

Proof of Proposition 7. Let $s^* = (0, y)$ be a boundary limit. The Jacobian for the dynamical system is again

$$J = \begin{pmatrix} \frac{\lambda_{L,1}'(0)e_2(0,y) - \lambda_{L,1}(0)\partial_0 e_2(0,y)}{\lambda_{L,1}(0)^2} & -\frac{\partial_y e_2(0,y)}{\lambda_{L,1}(0)} \\ -\frac{\partial_0 e_1(0,y)}{\lambda_{L,2}(y)} & \frac{\lambda_{L,2}'(y)e_1(0,y) - \lambda_{L,2}(y)\partial_y e_1(0,y)}{\lambda_{L,2}(y)^2} \end{pmatrix}.$$

We have $e_1(s^*) = 0$ and $\lim_{s_i \to 0} \lambda_{1,L}(s_i) = \infty$. Thus, the Jacobian becomes

$$J = \begin{pmatrix} \frac{\lambda'_{L,1}(0)e_2(0,y)}{\lambda_{L,1}(0)^2} & 0 \\ -\frac{\partial_0 e_1(0,y)}{\lambda_{L,2}(y)} & -\frac{\partial_y e_1(0,y)}{\lambda_{L,2}(y)} \end{pmatrix}.$$

From condition (1), it follows that there is an $a > 0$ such that $\lambda'_{L,1}(0)/\lambda_{L,1}(0)^2 = a$. We now substitute e_i for each $i = 1, 2$,

$$J = \begin{pmatrix} -a\left(\alpha(s^*)h_2(s_2)H - 1\right) & 0 \\ -\frac{\partial_{s_1}\alpha(s^*)}{\lambda_{L,2}(s_2)}H & -\frac{\partial_{s_2}\alpha(s^*)}{\lambda_{L,2}(s_2)}H \end{pmatrix}.$$

It is easy to see that the associated Eigenvalues are $\rho_1 = -a(\alpha(s^*)h_2(s_2)H - 1)$ and $\rho_2 = -\partial_{s_2} H\alpha(s^*)/\lambda_{L,2}(s_2)$. Now, $e_1(s^*) = 0$ implies $\alpha(s^*)H - 1 = 0$. Thus, $\alpha(0, s_2)h_2(s_2)H > 1$ which implies $\rho_1 < 0$. Moreover, $\rho_2 < 0$, because α is increasing in each argument and $\lambda_{L,2}(s_2) > 0$. Thus s is asymptotically stable. □

Proof of Proposition 8. The necessity of part (i), (iii) and (v) follow from Proposition 2. We show that if $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a strategy profile which has $\phi = (\phi_1, \ldots, \phi_n)$ as its generalized inverse, where ϕ , satisfies part (ii) and (iv), then $\sigma_i(s_i) \in \arg\max_t u_i(t, s_i)$ for all $s_i \in [0, 1]$ and $i \in N$.

Fix s_i and set $t = \sigma_i(s_i)$. We consider the possible deviations $t' > t$ and $t' < t$ separately.

(1.) Suppose agent i with signal $s_i = \phi_i(t)$ chooses a stopping time

$t' > t$. Then

$$\frac{du_i(t', s_i)}{dt} = -\Pr(H|s_i)(1 - G_{i,H}(t'))rH + \Pr(L|s_i)\left(G'_{i,L}(t') + r(1 - G_{i,L}(t'))\right)$$
$$= -\Pr(L|s_i)r(1 - G_{i,L}(t'))\left(\frac{\Pr(H|s_i)}{\Pr(L|s_i)}\frac{1 - G_{i,H}(t')}{1 - G_{i,L}(t')}H - 1 + \frac{1}{r}\frac{G'_{i,L}(t')}{1 - G_{i,L}(t')}\right)$$
$$\leq -\Pr(L|s_i)r(1 - G_{i,L}(t'))\left(\frac{\Pr(H|s_i)}{\Pr(L|s_i)} - \frac{\Pr(H|\phi_i(t'))}{\Pr(L|\phi_i(t'))}\right)\frac{1 - G_{i,H}(t')}{1 - G_{i,L}(t')}H$$
$$\leq 0$$

where the third inequality follows by substituting (9) evaluated at t', noting that ϕ_i is decreasing by hypothesis, and

$$\frac{\Pr(H|s_i)}{\Pr(L|s_i)} = \frac{p_0}{1 - p_0}\frac{F'_{i,H}(s_i)}{F'_{i,L}(s_i)}$$

is increasing by MLRP. Thus, $t' > t$ cannot be optimal.

(2.) Suppose agent i with signal $s_i = \phi_i(t)$ chooses a stopping time $t' < t$. First, consider t' in the image of σ_i. Using (9) we obtain

$$\frac{du_i(t', s_i)}{dt} = -\Pr(H|s_i)(1 - G_{i,H}(t'))rH + \Pr(L|s_i)\left(G'_{i,L}(t') + r(1 - G_{i,L}(t'))\right)$$
$$= -\Pr(L|s_i)r(1 - G_{i,L}(t'))\left(\frac{\Pr(H|s_i)}{\Pr(L|s_i)}\frac{1 - G_{i,H}(t')}{1 - G_{i,L}(t')}H - 1 + \frac{1}{r}\frac{G'_{i,L}(t')}{1 - G_{i,L}(t')}\right)$$
$$= -\Pr(L|s_i)r(1 - G_{i,L}(t'))\left(\frac{\Pr(H|s_i)}{\Pr(L|s_i)} - \frac{\Pr(H|\phi_i(t'))}{\Pr(L|\phi_i(t'))}\right)\frac{1 - G_{i,H}(t')}{1 - G_{i,L}(t')}H$$

where in contrast to (1.), the third line is now an equality. Since $\phi_i(t') > s_i$, it follows that $du_i(t', s_i)/dt > 0$, so that t' cannot be optimal. Second, consider $t' < t$ outside the image of σ_i. Because t' is outside the image of σ_i, by (9),there exists an agent i with signal $s'_i > s_i$ who prefers to stop at a time $t'' > t'$. Therefore, by monotonicity of agent i's best response (see Lemma 2), agent i with signal s_i, also prefers to stop at t''. Since this argument applies to all t' outside the image of σ_i, such a deviation is never optimal. □

Proof of Proposition 9. Part (i): At an interior limit, we have $Y_i(s^*) = 0$ for all $i \in A$. This is possible only if there exists an $u^* > 1$ such that $\tilde{u}_i(s_i^*) =: u^*$ for all $i \in A$. Thus $Y_i(s^*) = 0$ is equivalent to $u^* = 0$.

Part (ii): By Aggregate Uncertainty, if there are two agents $i \neq j$ such that $s_i^* = s_j^* = 0$, then $\tilde{u}_i(s^*) = u_j^*(s^*) < 0$. But then there exists a finite time t such that i prefers not to stop after t which contradicts the hypothesis that s^* is a limit point with $s_i^* = 0$. Now, let $s_i^* = 0$. Then $s_j^* > 0$ and thus $\lambda_{j,L}(s_j^*) > 0$ for all $j \in A \setminus \{i\}$. Therefore, $Y_j(s^*) = 0$ for all $j \in A \setminus \{i\}$ is possible only if there exists u^* such that $u_j(s^*) = u^*$. Subsituting this back into $Y_j(s^*)$ for $i \neq j$, we obtain have

$$Y_j(s^*) = 0 \Leftrightarrow \frac{1}{|A|-1} \sum_{j \in A} \frac{\tilde{u}_j(s^*)}{1 - p_j(s^*)} - \frac{\tilde{u}_i(s^*)}{1 - p_i(s^*)} \quad = \frac{1}{|A|-1} \frac{\tilde{u}_i(s^*)}{1 - p_i(s^*)} = 0$$

Where the last equation implies $\tilde{u}_i(s^*) = 0$. Equation (2) then implies that $\tilde{u}_i(s_i^*) = 0$. Finally, again by aggregate uncertainty, we have $h_j(s_j^*) > 1$ for all $j \in A \setminus \{i\}$. Thus,

$$\alpha(s^*) h_j(s_j^*) H - 1 \geq \alpha(s^*) H - 1 = 0.$$

It then follows from Equation (2) that $u^* > 0$. □

Proof of Proposition 10. Part (i): *We show that if an interior limit does not exist, then some agent has a dominant signal.* Define

$$S_0 = \{(s_1, \ldots, s_n) |\, \tilde{u}_i((s_i, s_{-i}) \leq 0 \text{ for all } i \text{ with equality for some } i\}.$$

An interior limit exists if there is an $(s_1, \ldots, s_n) \in S_0$ and $i \neq j \in N$ such that $\tilde{u}_i(s_i, s_{-i}) = \tilde{u}_j(s_j, s_{-j}) = 0$. Because the signal distributions are differentiable and have full support, any curve in S_0 is continuous. Therefore, if there is no interior limit, there must be an agent i such that $\tilde{u}_i(s_i, s_{-i}) > \tilde{u}_j(s_j, s_{-j}) = 0$.

Suppose agent i is the agent, for whom the last statement is true. It then follows that $\tilde{u}_i(0, (1, \ldots, 1)) < 0$. Otherwise, there would exists s with $s_i = 0$ and $s_{-i} < (1, \ldots, 1)$ such that $0 = \tilde{u}_i(0, s_{-i}) \leq \tilde{u}_j(s_j, s_{-j})$ for every $j \neq i$, contradicting the hypothesis. Thus, there is an $\hat{s}$ with $\hat{s}_i > 0$ and $\hat{s}_{-i} = (1, \ldots, 1)$ such that $0 = \tilde{u}_i(\hat{s}_i, \hat{s}_{-i}) > \tilde{u}_j(\hat{s}_j, \hat{s}_{-j})$. The last inequality implies that agent i has a dominant signal.

Part (ii): The proof is easy. The details are omitted. □

References

Aidt, T. S. and P. S. Jensen (2009). The taxman tools up: An event history study of the introduction of the personal income tax. *Journal of Public Economics 93*(1-2), 160–175.

Banerjee, A. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics 107*(3), 797–817.

Becker, M. H. (1970). Factors affecting diffusion of innovations among health professionals. *American Journal of Public Health and the Nations Health 60*(2).

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy 100*(5), 992–1026.

Bliss, C. and B. Nalebuff (1984). Dragon-slaying and ballroom dancing: The private supply of a public good. *Journal of Public Economics 25*, 1–12.

Bolton, P. and C. Harris (1999, mar). Strategic Experimentation. *Econometrica 67*(2), 349–374.

Chamley, C. and D. Gale (1994). Information Revelation and Strategic Delay in a Model of Investment. *Econometrica 62*(5), 1065–85.

Décamps, J.-P. and T. Mariotti (2004, sep). Investment timing and learning externalities. *Journal of Economic Theory 118*(1), 80–102.

Eeckhoudt, L. and C. Gollier (1995). Demand for risky assets and the monotone probability ratio order. *Journal of Risk and Uncertainty 11*, 113–122.

Foster, A. D. and M. R. Rosenzweig (2010). Microeconomics of Technology Adoption. *Annual Review of Economics 2*, 395–424.

Fudenberg, D., D. M. Kreps, and D. K. Levine (1988, apr). On the robustness of equilibrium refinements. *Journal of Economic Theory 44*(2), 354–380.

Fudenberg, D. and J. Tirole (1991). *Game theory*. MIT Press.

Goolsbee, A., T. Holmes, B. Jovanovic, A. Kashyap, L. Katz, S. Levitt, and P. P. J. Klenow (2002). Evidence on Learning and Network Externalities in the Diffusion of Home Computers. *The Journal of Law and Economics 45*(2), 317–343.

Hofbauer, J. and K. Sigmund (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

Kaustia, M. and S. Knüpfer (2012, may). Peer performance and stock market entry. *Journal of Financial Economics 104*(2), 321–338.

Keller, G., S. Rady, and M. Cripps (2005). Strategic Experimentation with Exponential Bandits. *Econometrica 73*(1), 39–68.

Liu, H., Q. Sun, and Z. Zhao (2014, jan). Social learning and health insurance enrollment: Evidence from China's New Cooperative Medical Scheme. *Journal of Economic Behavior and Organization 97*, 84–102.

Malueg, D. and S. Tsutsui (1997). Dynamic R&D competition with learning. *The RAND Journal of Economics 28*(4), 751–772.

Moscarini, G. and F. Squintani (2010). Competitive experimentation with private information: The survivor's curse. *Journal of Economic Theory 145*(2), 639–660.

Murto, P. and J. Välimäki (2011, may). Learning and Information Aggregation in an Exit Game. *The Review of Economic Studies 78*(4), 1426–1461.

Olsen, M. (1965). *The Logic of Collective Action*. Harvard Economic Studies. Harvard University Press.

Palfrey, T. R. and H. Rosenthal (1984). Participation and the provision of discrete public goods: a strategic analysis. *Journal of Public Economics 24*(2), 171–193.

Rosenberg, D., E. Solan, and N. Vieille (2007, nov). Social Learning in One-Arm Bandit Problems. *Econometrica 75*(6), 1591–1611.