

This is a repository copy of *Who goes first? Strategic delay under information asymmetry*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/142614/>

Version: Published Version

Article:

Achim, Peter (2018) *Who goes first? Strategic delay under information asymmetry*. *Theoretical Economics*. pp. 341-375. ISSN 1555-7561

<https://doi.org/10.3982/TE2171>

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial (CC BY-NC) licence. This licence allows you to remix, tweak, and build upon this work non-commercially, and any new works must also acknowledge the authors and be non-commercial. You don't have to license any derivative works on the same terms. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Who goes first? Strategic delay under information asymmetry

PETER A. WAGNER

Department of Economics, University of Bonn

This paper considers a timing game in which heterogeneously 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 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, there can be a persistent wedge in posterior beliefs between well and poorly informed agents, because the poorly informed tend to “drive out” the well informed.

KEYWORDS. Strategic delay, social learning.

JEL CLASSIFICATION. D83, C73.

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 into 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 Forssmann is of someone who took action in an environment of “wait and see,” in which everyone hopes for the independent initiative of a volunteer who resolves some of the risks related 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, that fought for commercial access to AT&T’s telephone networks in the 1960s, facing substantial

Peter A. Wagner: paw@uni-bonn.de

I am indebted to Martin Osborne, Colin Stewart, Ettore Damiano, Pauli Murto, Sven Rady, and three anonymous referees, 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.

Copyright © 2018 The Author. Theoretical Economics. The Econometric Society. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at <http://econtheory.org>.
<https://doi.org/10.3982/TE2171>

legal fees and significant risk. The legal procedure ended with a favorable ruling by the 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 and Klenow 2002), stock market entry (Kaustia and Knüpfer 2012), and the introduction of 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 generates a second-mover advantage that provides agents with an incentive to free-ride on others' initiatives.

I characterize the Bayes–Nash equilibria as allowing for a heterogeneous distribution of information. The equilibria can be broadly classified into two types. Equilibria may end with some agent's immediate investment if the information structure is capable of generating an informational cascade. If the information structure prohibits cascades, 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 beliefs and outcomes. 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 increasingly less likely to ever make the investment. 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, because the poorly informed tend to "drive out" the well informed.

Intuitively, this effect is the result of an informational feedback loop. Note that for an agent with an accurate signal about the state, the highest possible posterior belief about the state is higher than that of a poorly informed agent. Similar to the logic for equilibria in mixed strategies, in an equilibrium with delay, the poorly informed agent stops more

quickly than the agent with superior information, because the most optimistic types must be given an incentive to delay their investment. While the presence of such optimistic types depends on the true state, the behavior of the uninformed depends only on his *own* information. Stopping with a higher probability, the poorly informed reveals more information through inaction and thus continues to be worse informed. This reinforces a cycle in which one agent continues to be less informed than others while having to stop at a higher rate. In the limit, the less informed agent eventually invests if no other agent does, while the better informed are less and less likely to invest, despite the fact that their expected return remains positive and bounded away from zero.

In this type of equilibrium, some agent invests with certainty regardless of the state of the world, which might be in excess of the social optimum. The cost of this excess is borne mainly by the poorly informed. 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, 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 \(1984\)](#). 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 that allow 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 individual cost of contributing to a public good of commonly known value. In contrast, I assume that agents are privately informed about the uncertain return to an investment. The public good is the information about the true return, and the cost of contributing is the loss incurred when receiving a negative return. Thus, there is not only the strategic interaction that arises from the presence of asymmetric information, but also learning about the state of the world from observing others’ behavior.

Observational learning has been studied by a substantial number of papers, following the seminal articles of [Bikhchandani et al. \(1992\)](#) and [Banerjee \(1992\)](#). These papers consider models in which agents learn from other agents’ actions about a common state. [Bikhchandani et al. \(1992\)](#) show that when actions are a coarse signal about the state, then sequential decision making can lead to informational cascades in which agents ignore their own information and herd on an action that may be socially undesirable. They assume that agents choose their actions in an exogenously determined sequence, but this is not an essential requirement. [Chamley and Gale \(1994\)](#) find that cascades persist even with endogenous timing when agents can respond sufficiently quickly to

others' actions. The reason is that short delays can reveal only small amounts of information that may not suffice to change an agent's decision from one action to another if the decision he has to make is binary. This explains why informational cascades do not arise when agents have to choose an action from a continuum, as in [Gul and Lundholm \(1995\)](#) and [Murto and Välimäki \(2013\)](#).

The essential novelty in this paper compared to the social learning literature is that agents not only learn from the behavior of others, but also from the experience of others. This additional source of information increases the agents' option value of delay and can thus induce them to wait longer. As a result, delay may arise naturally in equilibrium. Delay and information cascades are in fact mutually exclusive. That is, delay occurs in a robust equilibrium if and only if there is no information cascade. This finding is consistent with [Chamley and Gale \(1994\)](#), in which the game ends immediately in the continuous time limit when agents learn only from others' actions, as well as with [Gul and Lundholm \(1995\)](#) and [Murto and Välimäki \(2013\)](#), in which equilibrium exhibits delay and no cascades arise.

The effects of pure informational spillovers from payoff observability has been studied in the strategic experimentation literature starting with [Bolton and Harris \(1999\)](#) and [Keller et al. \(2005\)](#). In these papers, 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 and thus there is an incentive for agents to free-ride on other agents' experimentation. This in turn leads to inefficient levels of experimentation in equilibrium. One can view the present model as a limit of these experimentation games, in which choosing the risky action is immediately fully revealing. Under symmetric information, the limit game is then isomorphic to the public goods model of [Bliss and Nalebuff \(1984\)](#).

A number of papers study versions of games of strategic experimentation with asymmetrically informed agents. Some of these papers study non-competitive models in which agents are privately informed about their cost of delay. In [Décamps and Mariotti \(2004\)](#), agents learn over time about a common state variable by observing a public signal and from the experience of others. Agents have an incentive to delay investment to signal high cost; thus actions do not reveal information about the state of the world, as is the case here.

[Rosenberg et al. \(2007\)](#) and [Murto and Välimäki \(2011\)](#) study a model of strategic experimentation with private payoffs and public exits. Agents are ex ante symmetric and information about the common state is accumulated through experimentation over time. Public exits reveal information about the state, but agents do not learn about the state from the experience of their peers. Therefore, there are no free-riding incentives. Moreover, in all three papers, agents are ex ante symmetric. The focus of the present paper is to explore the effects of asymmetries in the distribution of information.

Another array of papers considers a 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](#)). In these papers, investment is competitive due to patent restrictions, so that there is no free-riding motive to delaying investments. [Daron et al. \(2011\)](#) consider a patent race with privately informed agents

in which free-riding emerges due to limited patent protection. The authors choose the information structure specifically to eliminate any free-riding motive.

The remainder of the paper is structured as follows. Section 2 contains the model and assumptions, Section 3 covers efficiency, Section 4 contains the equilibrium analysis that is discussed in Section 5. Section 6 concludes.

2. MODEL

There is a set of agents $N = \{1, \dots, n\}$ who face the option to invest in identical projects that yield an uncertain return that depends on the realization of an unknown state of the world $\theta \in \{H, L\}$, where $H > 0$ 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, in the order of their indices, whether to invest immediately. Preemption allows agents to move sequentially at time zero without delay, which is essential for equilibrium existence and for establishing an appropriate efficiency benchmark.¹ When no agent preempts, the game enters the *waiting phase* in which each agent delays his investment by a positive amount of time. We 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 agent i 's stopping time conditional on reaching t_i in the waiting phase. When $t_i = \infty$, agent i waits indefinitely.

The first agent to make the investment immediately realizes the return and thereby reveals the state to everyone. Let $t_{-i} = (t_j)_{j \neq i}$ denote the profile of stopping times of all agents other than i . The payoff for agent i from stopping at t_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 with a smooth cumulative distribution function $F_{i, \theta}(\cdot)$, which we assume has full support and a bounded density. A strategy for agent i is a function $\sigma_i : [0, 1] \rightarrow [0, \infty] \cup \{-i\}$ 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 $s_i \in [0, 1]$.

For a given strategy profile $(\sigma_i)_{i \in N}$, let $\tau(s) = \min_{i \in N} \sigma_i(s_i)$ be the 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, define

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

¹Without preemption, equilibria may fail to exist when some agent stops at time zero 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.

to be the set of agents who 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), which says that the likelihood ratio $F'_{i,H}(s_i)/F'_{i,L}(s_i)$ is increasing in s_i for each agent i . We 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 any other agents’ signals.

ASSUMPTION 2 (Aggregate uncertainty). *We have $E[\theta|s_i=0, s_j=0] < 0$ for any $i \neq j$.*

The assumption of aggregate uncertainty is important to focus the analysis on informational rather than coordination problems. It says that if any pair of agents would commonly learn that both received their worst possible signal, then neither of them would be willing to make the investment. Aggregate uncertainty ensures that others’ private information does not only influence an agent’s timing of investment, as would be the case without this assumption, but also if an agent will choose to invest 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 $p_i(s_i, s_{-i})$ 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

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

Throughout, we call an agent i *more optimistic* than an agent j if agent i ’s posterior belief at a given public history is higher than agent j ’s. Moreover, we refer to agent i as being

better informed than agent j if the variance of agent i 's posterior belief is higher than the variance of agent j 's belief.

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

$$\tilde{u}_i(s_i, s_{-i}) := E[\theta | s_i, s_{-i}] = p_i(s_i, s_{-i})H - (1 - p_i(s_i, s_{-i})).$$

In some cases, agents may be endowed with particularly informative signals that dominate 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 $j \neq i$.

A dominant signal for agent i is a signal such that knowing that agent i 's expected stopping value is negative discourages even the most optimistic competitor.

We denote by $\alpha(s_1, \dots, s_n)$ 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

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

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

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

Denote by h_i the *reverse hazard ratio* (RHR) for agent i at $s_i \in [0, 1]$, 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 ratio h_i and the likelihood ratio of the public posterior α allow us to decompose the public posterior belief about the state into the common component and a private component:

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

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 this decomposition, we write

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

The left-hand side measures the relative payoff from investment. The right-hand side shows that this value differs across agents only through differences in their respective RHRs.

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

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

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

This assumption is a regularity condition that ensures that the distribution of signals in a low state is well behaved around zero for all agents. Intuitively, it says that close to zero, the curvature of the distribution function for signals in a low state is neither too small nor too large relative to its slope. It is always satisfied, for example, if $F_{i,L}$ is concave with finite second derivative. 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 distributions of signals across states.

3. SOCIALLY OPTIMAL STOPPING

In this section, we introduce a notion of efficiency that addresses the question of how agents should behave so as 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 $E = E_1 \times \dots \times E_n$ 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

payoffs provided stopping is indeed socially desirable for all profiles in E . Finding the efficient strategy profile thus means determining the stopping region E that maximizes the expected welfare $\mathbb{E}[v(s) \mid s \in E]$, 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 every 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 $\hat{s} = (\hat{s}_1, \dots, \hat{s}_n) \in [0, 1]^n$ such that $\hat{\sigma}_i(s_i) = -i$ if $s_i \geq \hat{s}_i$ and $\hat{\sigma}_i(s_i) = \infty$ otherwise. The threshold profile satisfies $\tilde{v}_i(\hat{s}) \leq 0$ for all i with $\tilde{v}_i(\hat{s}) = 0$ whenever $\hat{s}_i < 1$, where*

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

Moreover, an efficient strategy profile exists.

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 nonpositive.

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 ratios are constants given by $h_1(s_1) = \beta$ and $h_2(s_2) = \gamma$, respectively. By the same logic as in (2), we have the inequality

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




FIGURE 1. Efficient stopping with two agents.

Since reverse hazard 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 is 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 that 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.

4.1 Preliminaries

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

PROPOSITION 2. *Let $(\sigma_1, \dots, \sigma_n)$ be a Bayes–Nash equilibrium. Then for each $i = 1, \dots, n$, we have the following scenarios:*

- (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 (countable) set of discontinuities of σ_i . Then σ_i is differentiable on $(s_i^-, s_i^+) \setminus D_i$.
- (iii) We have $|A(t)| \neq 1$ for a.a. $t \geq 0$.

Intuitively, the proposition says that each agent's equilibrium strategy is a decreasing function that has flat regions only at the upper and lower tails, 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, which, if nonempty, must contain at least two agents except on a set of measure zero.

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 from t until t' is equal to the expected loss avoided if another agent stops between these times and the state turns out to be low. However, the agent also incurs a loss from waiting for the case in which no agent stops due to discounting. Now, the higher an agent's signal, the higher his belief that the state is indeed high, which reduces the value of delaying investment. Thus, if it is a best response for an agent with signal s_i to stop at t , then no agent with signal $s'_i > s_i$ will stop later than t .

Equilibrium strategies are "smooth" in the articulated sense because payoffs are differentiable with respect to stopping times, and because the signal distributions are "well behaved" in the sense that they have full support with bounded, continuously differentiable densities. Therefore, a small variation in signals leads 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. One 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 the other imitates. We call this scenario *informed preemption*. The second possibility is that a strongly optimistic agent preempts the game regardless of the realization of his signal, while the other waits for him to move. We refer to this second scenario as *uninformed preemption*.

Informed preemption In an equilibrium with informed preemption, one agent preempts the game if he expects his payoff to be nonnegative, and otherwise waits indefinitely; the other waits forever for sure. Formally, suppose agent i has a dominant signal, and let s_i^* be the signal solving $E[\theta|s_i^*] = 0$. Then, in an equilibrium with informed preemption, agent i preempts if $s_i \geq s_i^*$ and waits indefinitely otherwise, while the other agent waits indefinitely for sure. Informed preemption is possible in equilibrium if the preempting agent'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 a dominant signal so that agent i 's inaction at the beginning of the game conveys sufficiently bad news to the other to "overpower" any potential good news he might have himself. That the equilibrium conditions are satisfied follows immediately from the definition of dominant signals. Agent i 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.

The notion of dominant messages provides a necessary and sufficient condition for the existence of an equilibrium in which an information cascade arises with positive probability. It generalizes a finding in [Herrera and Hörner \(2013\)](#) in which information cascades form if reverse hazard ratios are monotonically decreasing. Here, if agents have identical signal distributions with a decreasing RHR, then each agent has a dominant signal. To see this, recall that by (1) we have

$$\frac{\tilde{u}(s, 1)}{1 - p(s, 1)} = \alpha(s, 1)h(s)H - 1 > \alpha(1, s)h(1)H - 1 = \frac{\tilde{u}(1, s)}{1 - p(1, s)},$$

where the second inequality holds when h is decreasing (we drop subscript i). Now, if $u(s, 1) = 0$, then $u(1, s) < 0$, which shows that each agent has a dominant signal.

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 that the other agent waits indefinitely, and waiting indefinitely is a best response to the first 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 agent supposed to respond? In equilibrium, agent 2 waits indefinitely after such a deviation of agent 1, regardless of his own belief.

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 with uninformed preemption is sustained by pure action. In contrast, in equilibria with informed preemption, inaction generates an information cascade that makes the equilibrium robust to small perturbations in payoffs.

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 fixed points of the associated dynamical system. Fixed 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 fixed points and their stability attributes affect equilibrium properties.

When there are only two agents in the game, it follows from Proposition 2 that for an equilibrium with delay (σ_1, σ_2) , each strategy σ_i is differentiable at a.a. $s_i < s_i^+$. Moreover, each agent’s strategy has a differentiable monotone inverse, which we can use to derive a system of differential equations whose solutions are candidates for inverse equilibrium strategies.

Fixing a pair of inverse equilibrium strategies (ϕ_1, ϕ_2) , monotonicity implies that we can write the probability that agent i stops before time t in state θ as $1 - F_{i,\theta}(\phi_i(t))$ for each $i = 1, 2$. Therefore, the expected payoff from stopping at time $t > 0$ is given by

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

The first and second terms 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 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

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

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

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

The left-hand side is agent i 's marginal cost of waiting, the right-hand side is the expected marginal gain from delaying investment by another instant. With probability $\lambda_{-i,L}(\phi_{-i}(t))\phi'_{-i}(t)$, the other agent stops at that instant, and then agent i avoids the loss of -1 if the state is low, which is the case with probability $1 - p_i(\phi_i(t), \phi_{-i}(t))$.

Now, for any equilibrium (σ_1, σ_2) with delay, the pair of inverses (ϕ_1, ϕ_2) must solve the system of differential equations

$$\begin{aligned} -\phi'_1(t) &= Y_1(\phi_1(t), \phi_2(t)), \\ -\phi'_2(t) &= Y_2(\phi_1(t), \phi_2(t)), \end{aligned} \tag{3}$$

where

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

By **Proposition 2**, strategies belonging to an equilibrium with delay must be monotonically decreasing, so that a solution path can belong to an equilibrium if and only if it is strictly decreasing. The following proposition shows that monotonicity and differentiability are in fact sufficient conditions for equilibrium.

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, differentiable inverse strategies solving (3) 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 fixed points of the dynamical system (3) 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 that 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 eventually converges to one of the fixed 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.

Interior limits An interior limit is a fixed point (s_1, s_2) of (3) 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 (2), the point (s_1, s_2) must lie at the intersection of the zero-profit curves.

The following proposition shows that an interior limit exists if neither has or both agents have a dominant signal, and that any interior limit is also a limit of an equilibrium with delay.

PROPOSITION 6. *There exists an equilibrium that converges to an interior limit $s^* = (s_1^*, s_2^*)$ if both have or neither agent has a dominant signal. Moreover, if $\varphi'_i(s_i^*)\varphi'_{-i}(s_{-i}^*) < 1$, then this equilibrium is unique.*

When no agent has a dominant signal, then, letting $\hat{s}_i = \varphi_i^{-1}(1)$, 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 upward ($\tilde{u}_1 < 0$) 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 downward ($\tilde{u}_1 > 0$) and to the right ($\tilde{u}_2 > 0$), thus moving toward the point of intersection.

The stability attributes of an interior limit determine 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 (3) is locally Lipschitz in the interior and thus for any initial interior point s , there exists a unique solution. Now, starting at a stable interior limit s^* , we can choose any $s > s^*$ in a small neighborhood around the fixed point, and solve (3) backward 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 fixed 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 agent it must

²When $\lim_{s_i \rightarrow 1} F'_{i,L}(s_i) = 0$.

remain positive.³ This implies that at a 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 \hat{s}_i , such that for any $s_i^\dagger \leq \hat{s}_i$, 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 1 is zero if he observes his worst signal and learns that agent 2 observed a signal no higher than s_2 . It is easy to check that $s^* = (0, s_2)$ is a boundary limit of the dynamical system (3). Notice that $\tilde{u}_1(0, s_2) < \tilde{u}_2(s_2, 0)$, since agent 1 knows only that s_2 is an upper bound of agent 2’s signal.

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. Once we have determined that a boundary point is stable, it follows from Lipschitz continuity (shown in the proof to [Proposition 2](#)) that there exists a unique solution path that ends at s^* . 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, there exists a unique interior limit. The left panel shows the phase diagram for a 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 a 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 (possibly involving preemption with positive probability) that converge to the unique interior limit. Decreasing RHR implies that neither agent is strongly optimistic, and thus there is no

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

⁴Note 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).

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. In such large 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' signals.

A substantial difference in larger games arises in equilibria with delay. When there are more than two agents in the game, each agent can become a passive bystander during the course of the waiting phase. Recall that with two agents, 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 the source of this complexity, note that 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 0. 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 (ϕ_1, \dots, ϕ_n) of inverse strategies constitutes an equilibrium if there exists an equilibrium $(\sigma_1, \dots, \sigma_n)$ 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 $F_{i,\theta}(\phi_i(t))$. By conditional independence, the probability that at least one agent other than i stops before time t is given by

$$G_{i,\theta}(t) := 1 - \prod_j F_{j,\theta}(\phi_j(t)). \tag{4}$$

Since ϕ_i is 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. The expected payoff for agent i from choosing stopping time t is

$$\Pr(H|s_i) \left(\int_0^t e^{-r\tau-i} dG_{i,H}(\tau-i) + (1 - G_{i,H}(t))e^{-rt} \right) H + \Pr(H|s_i)G_{i,H}(0)H - \Pr(L|s_i)(1 - G_{i,L}(t))e^{-rt}.$$

The interpretation is analogous to the two-agent case. The derivation of the equilibrium conditions follows essentially the same steps as before. A sufficient condition for agent i to be willing to delay stopping is that his marginal expected payoff is greater than zero:

$$- \Pr(H|\phi_i(t))(1 - G_{i,H}(t))rH + \Pr(L|\phi_i(t))(G'_{i,L}(t) + r(1 - G_{i,L}(t))) \geq 0. \tag{5}$$

Now, substitute the stopping distribution $G_{i,\theta}$ from (4) as well as

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

divide both sides of (5) by the total probability of reaching time t , and substitute agent i ’s posterior belief p_i to obtain the condition

$$r\tilde{u}_i(\phi_i(t), \phi_{-i}(t)) \leq -(1 - p_i(\phi_i(t), \phi_{-i}(t))) \sum_{j \neq i} \lambda_{j,L}(\phi_j(t)) \phi'_j(t). \tag{6}$$

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

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

PROPOSITION 8. *A profile (ϕ_1, \dots, ϕ_n) of inverse strategies constitutes an equilibrium if the following conditions hold:*

- (i) *Every ϕ_i is continuous, differentiable a.e., and weakly decreasing.*
- (ii) *The inequality $\phi_i(-i) < 1$ implies $\phi_j(-j) = 1$ for all $j \neq i$.*
- (iii) *For every $i \in N$, condition (6) holds with equality at a.a. $t > 0$.*

Continuity follows from strict monotonicity on the support of types that stop with delay in finite time. Part (ii) says that at most one agent preempts the game. Part (iii) says that for any active agent, the strategy is pinned down by first-order conditions. Note that there is no restriction on the subset of active agents.

Take as given a profile (ϕ_1, \dots, ϕ_n) that satisfies the conditions in Proposition 8. Using elementary operations and rearranging the equation system obtained by setting (6) equal for each $i \in A(t)$, we isolate the derivatives of inverse strategies of active agents. This yields the dynamical system

$$\begin{aligned}
 -\phi'_1(t) &= \mathbf{1}_{\{1 \in A(t)\}} \cdot Y_1(\phi_1(t), \dots, \phi_n(t)) \\
 -\phi'_2(t) &= \mathbf{1}_{\{2 \in A(t)\}} \cdot Y_2(\phi_1(t), \dots, \phi_n(t)) \\
 &\vdots \\
 -\phi'_n(t) &= \mathbf{1}_{\{n \in A(t)\}} \cdot Y_n(\phi_1(t), \dots, \phi_n(t)),
 \end{aligned}
 \tag{7}$$

where

$$Y_i(s_1, \dots, s_n) = \frac{r}{\lambda_{i,L}(s_i)} \left(\frac{1}{|A(t)| - 1} \sum_{j \in A(t)} \frac{\tilde{u}_j(s_j, s_{-j})}{1 - p_j(s_j, s_{-j})} - \frac{\tilde{u}_i(s_i, s_{-i})}{1 - p_i(s_i, s_{-i})} \right)
 \tag{8}$$

and $\mathbf{1}_{\{i \in A(t)\}}$ is an indicator function that takes the value 1 if agent i is active and 0 otherwise.

In contrast to the two-agent case, some agents may become passive in the limit, that is, some agents may not invest with positive probability beyond some finite time. The specification of equilibrium limit points is, therefore, more delicate than in the two-agent case because we must account for those agents who become passive in the limit.

We define an equilibrium limit to be a profile $s^* = (s_1^*, \dots, s_n^*)$ such that there is a set $A \subseteq N$ with $|A| \geq 2$, so that

$$\begin{aligned}
 Y_i(s_1^*, \dots, s_n^*) &= 0 \quad \forall i \in A, \\
 \tilde{u}_i(s_i^*, s_{-i}^*) &< \max_{j \in A} \tilde{u}_j(s_j^*, s_{-j}^*) \quad \forall i \notin A.
 \end{aligned}
 \tag{9}$$

In other words, an equilibrium limit is a fixed point of the dynamical system (7), 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 that inactivity is the result of inferior information. It is easy to see that, for a given signal profile s and set A , condition (6) cannot hold along any solution path approaching s when the inequality is violated.

Analogous to the two-agent case, we call a limit point an interior limit if $s_i^* > 0$ for all i and a boundary limit otherwise. In the following proposition, we provide a characterization of the stopping values at these limits.

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

- (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 in the sense that for every agent who is active in the limit, the stopping value is zero. 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 whose limit signal lies on the boundary has a stopping value of zero. The stopping value of all other active agents equalizes as in the interior limit case, but it remains positive. This result is easy to see for the case of two agents, but the proposition generalizes to any number of participants in the game.

To see that the stopping value of all agents must be zero at any interior limit s^* , notice that for (8) to be equal to zero, the expression in parentheses 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 assumption 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 $\tilde{u}_j(s_j^*, s_{-j}^*) = 0$. 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 statements hold:*

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

⁵We 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 receive a positive stopping value.

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

If no interior limit exists, then it must be the case that one agent is more optimistic than all other agents. Thus, this agent must have a dominant signal. Part (i) implies that 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 beliefs of the remaining active agents either converge or, if convergence is impossible, all but one eventually become passive.

5. DISCUSSION

Equilibrium description

In general, we can classify equilibria by the timing of investments and by the dispersion of beliefs in the long run. Equilibria without delay exist when the information structure is distributed in such a way that one agent's decision to delay investment can prompt others to do the same. In this way, observational learning from inaction can lead to what Bikhchandani et al. (1992) calls an *informational cascade*, wherein an agent's action is independent of his private information.

Delay is unavoidable when the information structure does not admit informational cascades. In fact, all equilibria exhibit delay when no agent is strongly optimistic and none has a dominant signal. This is the case, for example, when highest signals are perfectly informative and the prior belief is low.⁶ In this case, no agent's inaction can with certainty dissuade all others from investing. Among equilibria with delay, there are two structurally different long-run outcomes. In the first type of equilibrium, a group consisting of the most optimistic agents remains active in the limit and their posterior beliefs converge over time. In the second type of equilibrium, posteriors do not converge; one of the agents, who is endowed with a high prior and a weak signal, is most likely to invest, despite the fact that some others are better informed about the state. Note, however, that in equilibria with delay, it is possible (and sometimes necessary) that there is an agent who preempts the game with positive probability.

In general, we find that in the context of indirect learning through delay, better information reduces the strategic incentive to invest first. That this is the case can be seen immediately by closer inspection of (7). An active agent who has very accurate private information, as measured by the divergence between his private belief and the public belief, stops at a lower rate than an agent who is endowed with less informative signals. More formally, better information is associated with a larger value of his reverse hazard ratio, and the larger this value, the lower the stopping rate.

Moreover, in the presence of a strongly optimistic agent, learning from others' delay tends to reinforce informational asymmetries over time. To see this in the case of two agents, reconsider the phase diagram shown in Figure 2. The left panel, shows the

⁶For example, consider the case in which $f_L(s_i) = 2s_i(1 - s_i)$ and $f_H(s_i) = s_i$ for each agent i . If $p_0 < 1/H$, then no equilibrium with preemption exists.

case of two strongly optimistic agents whose signals are drawn from identical distributions. There is an interior limit, but the limit is a saddle point, and thus the equilibrium converging to the interior limit is unique.

The instability of the interior limit is linked to the strategic effects of informational asymmetry on equilibrium dynamics. Along the equilibrium path approaching an interior limit, agents remain similarly well informed. However, after any displacement away from the equilibrium path, the best-response dynamics force the agent with relatively less information to reveal more and more information over time. In contrast, boundary limits are asymptotically stable. Therefore, when the information asymmetry at the outset is large enough, some agent reveals all information and invests with certainty in any equilibrium with delay.

We can summarize these results as follows. When there are large asymmetries in the distribution of information (such that there is no interior limit), there are two plausible outcome predictions. Either a well informed agent preempts the game and there is no delay or there is delay, in which case a less informed agent is most likely to invest first. In both cases, the asymmetry in the distribution of information shifts responsibility of providing information to one of the parties. When information is distributed more evenly, such that an interior point does exist, then there is an equilibrium in which beliefs converge over time, and the burden of investing first is distributed more equally.

Welfare properties

Free-riding incentives lead to inefficiencies in the timing of investments and in the extent of aggregate investment levels. In equilibria with delay, agents defer their investments to avoid their individual risk of suffering a loss and, instead, wait for others to provide this information for free. Delays serve no purpose in aggregating information, since any aggregate investment through decentralized stopping strategies can be implemented without delay in the preemption phase.

The agents' failure to internalize the social value of their own investment can have two, very different implications for aggregate investment levels in equilibrium. In equilibria with informed preemption, the preempting agent does not take into account the informational value of his decision for others. Thus, aggregate investment is below the socially optimum. Similarly, in an equilibrium with delay that converges to an interior limit, each active agent decides whether to invest based only on his individual expected return, ignoring the informational value for others. Therefore, aggregate investment is below the social optimum. Moreover, agents who are active in equilibrium might differ from those who are active in the efficient benchmark, which can create an additional source of inefficiency. The observation that free-riding incentives lead to suboptimally low aggregate investment levels is intuitive and a standard result in the related literature concerned with the private provision of public goods.

A novel insight is that in the presence of information asymmetry, free-riding incentives in conjunction with the option to delay investments can induce an informational feedback loop that results in excess aggregate investment. This dynamic arises in equilibria with delay that converge to a boundary limit. Along the equilibrium path of such

equilibria, the expected rates at which each agent stops must provide the most optimistic types among all others with an incentive to delay their investment. The more accurate an agent's information, the more optimistic his highest type and the more rapidly other, less informed agents have to stop. This in turn implies that the least informed agents reveal more information through inaction than those with better information—a cycle that reinforces itself over time.

The asymmetry in the distribution of information is a prerequisite for this phenomenon. The cost of the excess investment in aggregate is borne predominantly by poorly informed agents. While the presence of optimistic types depends on the true state, the behavior of the poorly informed agents depends only on their *own* noisy signals. Therefore, the likelihood that a poorly informed agent stops is relatively larger in a low state. Thus, as far as production of information is concerned, there is no “exploitation of the great by the small” (Olson 1965), but rather the opposite.

Equilibria with uninformed preemption are somewhat special in two regards. On the one hand, they are the only type of equilibrium in which no information is exchanged through delay (since the preempting agent stops immediately for sure). On the other hand, when investment is socially optimal for any realization of signal profiles (e.g., if n is large), then uninformed preemption is efficient. Intuitively, when the social gain from investing is large and independent of the agent's aggregate information, then it is efficient to rely on an uninformed agent to reveal the true state to everyone for sure. All other equilibria are, in contrast, inefficient.

Arrival of information

Questions relating to asymmetric information in timing games have been studied extensively in the literature, and it is worthwhile exploring the connections to and differences from the present paper in more detail. One set of papers considers timing games in which agents receive a private signal at the outset, and any additional information is obtained through inference from the timing of others' exits. Chamley and Gale (1994) consider a discrete time game in which agents are privately informed about their own presence, which is positively correlated with a binary state. Agents choose when to exit, and exits are publicly observable. When the period length goes to zero, then in the unique symmetric equilibrium, an analogue to information cascades of Bikhchandani et al. (1992) appears and the game ends immediately with probability 1. Intuitively, the game ends immediately because agents delay their investment only if the amount of information in each period is large enough to potentially affect their decision. Here, exits reveal not only an agent's private signal, but also the true state of the world. The informational value of exits is thus substantial, so that it makes sense that delays persist.

Chamley and Gale (1994) also show that there is a bias toward underinvestment, meaning that when the number of agents grows large, there can be insufficient investment in a high state, but never excess investment in a low state. Here, only equilibria with delay are sensitive to an increase in the number of agents. Assuming symmetric signal distributions, the effects on an increase in the number of agents on aggregate investment depends on the properties of the underlying distribution of signals. In general, when the number of agents increases, then the critical signal for each agent moves

closer to 1. The effect on aggregate investment is ambiguous and depends on the reverse hazard ratio (RHR). While the probability that an individual agent invests only in a high state increases, the total probability that there is at least one agent who receives such a signal in a low state increases as well. By [Proposition 9](#), each agent's stopping value at an interior limit must vanish, and thus, by condition (1), we have $\alpha(s^*, \dots, s^*)h(s^*)H = 1$. Thus, if the RHR increases in s^* , then the public posterior must decrease, so that the accuracy of aggregate decision making becomes worse as the number of agents grows large.

A number of other papers consider timing games in which agents are uninformed a priori and information arrives in the form of state-dependent flow payoffs gradually over time. [Rosenberg et al. \(2007\)](#) considers a game of private strategic experimentation with two players who observe their own payoff and whether the other agent has stopped. The continuous influx of information changes long-run outcomes dramatically, because given a sufficient amount of time, the agents eventually learn the true state. Thus, information cascades in which no player invests do not arise in this model.

[Murto and Välimäki \(2011\)](#) consider a related problem with private experimentation of many agents whose irreversible “exit” from experimentation is publicly observable. They show that in the symmetric equilibrium, information aggregates in exit waves that arrive randomly over time. As in [Rosenberg et al. \(2007\)](#), the continuous arrival of information implies that information cascades can arise only on investment, since in the long-run, agents eventually learn the true state. The symmetric equilibrium combines features of informed preemption and equilibria with delay and interior limit.

Replicator dynamics

There is a strong connection between the equilibrium learning dynamics of our model and replicator dynamics that are commonplace in evolutionary game theory and theoretical ecology. The standard replicator dynamics characterize the changes in composition of a population over time as a function of its payoff or “fitness” in relation to the population average. These dynamics are captured by coupled first-order differential equations.

In the present model, the dynamics that characterizes the composition of private information across agents over time exhibits qualitatively the same properties. The connection becomes most obvious in the case of two agents. The first-order condition obtained in this case yields differential equations that generate dynamics identical to the models of “competing species” (a special case of the 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 unstable fixed 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.

6. CONCLUSION

The objective of this paper was to reveal some of the mechanisms that govern strategic investments in environments in which information is dispersed and agents learn through observation from others’ actions and experience. I fully characterize equilibria for games with two agents, and long-run equilibrium outcomes for larger games. 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 appears 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 a 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.

APPENDIX: PROOFS

PROOF OF PROPOSITION 1. Because of the monotone likelihood ratio property, expected payoffs are nondecreasing 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 to stop at any higher signal. Thus, the stopping region is characterized by a profile \hat{s} of thresholds. The optimal threshold profile solves

$$\max_{(\hat{s}_1, \dots, \hat{s}_n)} p_0 \left(1 - \prod_{i=1}^n F_{i,H}(\hat{s}_i) \right) nH - (1 - p_0) \left(1 - \prod_{i=1}^n F_{i,L}(\hat{s}_i) \right). \tag{10}$$

The associated Lagrangian is

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

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

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

together with the Kuhn–Tucker conditions $\rho_i s_i = 0$ and $\mu_i (1 - s_i) = 0$, and $\rho_i, \mu_i \geq 0$ for all $i \in N$. Note that the left-hand side is equal to $\tilde{v}(\hat{s})$, so that the first-order condition

can also be written as $\tilde{v}(\hat{s}_i, \hat{s}_{-i}) = \rho_i - \mu_i$, where $(\hat{s}_i, \hat{s}_{-i})$ denotes the profile of signals for which agent i 's signal is \hat{s}_i and the remaining signals are given by \hat{s}_{-i} .

First note that setting $\hat{s}_i = 0$ is never optimal by aggregate uncertainty. If $\hat{s}_i \in (0, 1)$, then $\rho_i = \mu_i = 0$, so that $\tilde{v}(\hat{s}_i, \hat{s}_{-i}) = 0$. If $\hat{s}_i = 1$, then $\rho_i = 0$ and $\mu_i > 0$. Thus, the first-order condition implies that $\tilde{v}(\hat{s}_i, \hat{s}_{-i}) \leq 0$.

Existence follows from the extreme-value theorem. Since signal distributions are continuous, the objective in (10) is continuous. The set of signal thresholds $[0, 1]^n$ is compact; thus a solution exists. □

DEFINITION 4. The strategy σ_i has an atom at t if there exists an open set A of signals in $[0, 1]$ such that $\sigma_i(s_i) = t$ for a.a. $s_i \in A$.

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 zero.*

PROOF. (i) *At most one agent preempts the game.* Suppose there are two agents $i > j$ who preempt with positive probability. Then for each signal s_i with $\sigma_i(s_i) = -i$, agent i would do strictly better by stopping at time $t = \epsilon$ for $\epsilon > 0$ small, so that preemption is no best response for i .

(ii) *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 nonincreasing.*

PROOF. We show that equilibrium payoffs are submodular. Let $q(s_i) = \Pr(H|s_i)$. Denote by $G_{i,\theta}(t)$ the probability agent i assigns to the event that at least one other agent stops before t in state θ . The payoff of stopping at time $t \geq 0$ 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

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

We can now use that

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

and substitute

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

to obtain the inequality

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

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 nonincreasing in s_i . □

LEMMA 3. Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be an equilibrium and let $\phi = (\phi_1, \dots, \phi_n)$ be its generalized inverse, where $\phi_i(t) = \sup\{s_i | \sigma_i(s_i) \geq t\}$ for each $i = 1, \dots, n$. 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 $F_{i,L}$ has full support by hypothesis, it follows then that σ_i is differentiable almost everywhere.

By definition of U^* , it follows that

$$\begin{aligned} \Delta U_i^*(t, t', s_i) &= -\left(q(s_i) \int_0^t e^{-rz} dG_{i,H}(z) H + e^{-rt} u_i^*(t, s_i)\right) \\ &\quad + q(s_i) \int_0^{t'} e^{-rz} dG_{i,H}(z) H + e^{-rt'} u_i^*(t', s_i) \\ &= -e^{-rt} u_i^*(t, s_i) + e^{-rt'} u_i^*(t', s_i) + q(s_i) \int_t^{t'} e^{-rz} dG_{i,H}(z) H. \end{aligned}$$

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

$$e^{-rt} u_i^*(t, s_i) - e^{-rt'} u_i^*(t', s_i) \geq q(s_i) \int_t^{t'} e^{-rz} dG_{i,H}(z) H. \tag{11}$$

We further have

$$\int_t^{t'} e^{-rz} dG_{i,H}(z) \geq e^{-rt'} \int_t^{t'} dG_{i,H}(z) = e^{-rt'} (G_{i,H}(t') - G_{i,H}(t)). \tag{12}$$

Using a zero addition, we find that

$$\begin{aligned} & e^{-rt} u_i^*(t, s_i) - e^{-rt'} u_i^*(t', s_i) \\ &= e^{-rt} u_i^*(t, s_i) - e^{-rt'} u_i^*(t, s_i) + e^{-rt'} u_i^*(t, s_i) - e^{-rt'} u_i^*(t', s_i) \\ &= (e^{-rt} - e^{-rt'}) u_i^*(t, s_i) + e^{-rt'} q(s_i) (G_{i,H}(t') - G_{i,H}(t)) H \\ &\quad - e^{-rt'} (1 - q(s_i)) (G_{i,L}(t') - G_{i,L}(t)), \end{aligned}$$

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

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

Now, use (11) and (12) 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, $r(t' - t) \geq e^{-rt} - e^{-rt'}$. Altogether, it follows that

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

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 that $L^* = \max_{(t,t') \in I \times I} L(t, t')$ 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

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

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 nonincreasing, 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.

(ii) Follows from Lemma 3.

(iii) Suppose $A(t) = \{i\}$ on some open interval $(t_0, t_1) \subset \mathbb{R}_+$. But then there is a type $s_i \in \phi_i((t_0, t_1))$ in the set of types of agent i that stop during this interval who 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. □

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 is

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

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

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

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 at $t > 0$, 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. □

PROOF OF PROPOSITION 6. The main argument is given in the text. It remains to characterize the stability properties of the interior limit. 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

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

By the implicit function theorem, the null clines φ_1 and 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

$$\varphi'_1(x) = -\frac{\partial_x e_1(x, y)}{\partial_y e_1(x, y)}, \quad \varphi'_2(y) = -\frac{\partial_y e_2(x, y)}{\partial_x e_2(x, y)}.$$

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 ρ_1 and ρ_2 have opposite signs, which implies that the interior limit is a saddle point and hence is unstable. Thus, there exists a unique trajectory (the separatrix) that converges to s^* , and this trajectory constitutes an equilibrium path. \square

PROOF OF PROPOSITION 7. The main argument is given in the text. It remains to show that a boundary limit is asymptotically stable. Define again

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

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

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

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

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

From **Assumption 3**, 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 $i = 1, 2$:

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

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. \square

PROOF OF PROPOSITION 8. We show that if $\sigma = (\sigma_1, \dots, \sigma_n)$ is a strategy profile that has the generalized inverse $\phi = (\phi_1, \dots, \phi_n)$ that satisfies properties (i)–(iii), then $\sigma_i(s_i) \in \arg \max_t \tilde{u}_i(t, s_i)$ for all $s_i \in [0, 1]$ and $i \in N$. Fix s_i and set $t = \sigma_i(s_i)$.

Let $t > 0$. Suppose agent i with signal $s_i = \phi_i(t)$ chooses a stopping time $t' \neq t$. By a slight abuse of notation, we write $u_i(t, s_i)$ as the expected payoff of agent i when stopping at time t . Then

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

where the third equality follows by substituting (5) 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 $du_i(t', s_i)/dt > 0$ if $t' < t$ and $du_i(t', s_i)/dt < 0$ if $t' > t$, which implies that t is a best response for agent i with signal s_i .

Now let $t = -i$. By (iii), we have $du_i(\sigma_i(s_i), s_i)/dt = 0$ for all $s_i \in (s_i^-, s_i^+)$. Thus, $\lim_{s_i \uparrow s_i^+} du_i(\sigma_i(s_i), s_i)/dt = 0$, where $s_i^+ = \inf\{s_i | \sigma_i(s_i) = -i\}$. It follows that $du_i(0, s_i) \leq 0$ for $s_i \geq s_i^+$. Thus stopping immediately is a best response for $s_i \geq s_i^+$, since, by condition (ii), no other agent preempts with positive probability. \square

PROOF OF PROPOSITION 9. (i) At an interior limit, we have $Y_i(s^*) = 0$ for all $i \in A$. Since the point is interior, we have $s_i^* > 0$ and thus $\lambda_{i,L}(s_i^*) < \infty$ for all $i \in A$. It then follows from (8) that

$$\frac{\tilde{u}_i(s^*)}{(1 - p_i(s^*))} = \frac{\tilde{u}_j(s^*)}{(1 - p_j(s^*))} \quad \forall i, j \in A.$$

This in turn implies that $p_i(s^*) = p_j(s^*) =: p^*$ and $\tilde{u}_i(s^*) = \tilde{u}_j(s^*) =: u^*$ for all $i, j \in A$. It follows that

$$Y_j(s^*) = \left(\frac{1}{|A| - 1} \right) \sum_{j \in A} \frac{u^*}{1 - p^*} - \frac{u^*}{1 - p^*} = \left(\frac{1}{|A| - 1} \right) \frac{u^*}{1 - p^*}.$$

Hence, $Y_j(s^*) = 0$ implies that $\tilde{u}_i(s^*) = \tilde{u}_j(s^*) = u^* = 0$ for all $i, j \in A$.

(ii) By aggregate uncertainty, if there are two agents $i \neq j$ such that $s_i^* = s_j^* = 0$, then $\tilde{u}_i(s^*) = \tilde{u}_j(s^*) < 0$. 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 $\lambda_{j,L}(s_j^*) < \infty$ for all $j \in A \setminus \{i\}$. By the same argument as in part (i), it follows that there are p^* and u^* such that $p_j(s^*) =: p^*$ and $\tilde{u}_j(s^*) =: u^*$ for all

$j \in A \setminus \{i\}$. We have

$$\begin{aligned} Y_j(s^*) &= \frac{1}{|A| - 1} \sum_{k \in A} \frac{\tilde{u}_k(s^*)}{1 - p_k(s^*)} - \frac{\tilde{u}_j(s^*)}{1 - p_j(s^*)} \\ &= \frac{1}{|A| - 1} \left((|A| - 1) \frac{u^*}{1 - p^*} + \frac{\tilde{u}_i(s^*)}{1 - p_i(s^*)} \right) - \frac{u^*}{1 - p^*} \\ &= \frac{1}{|A| - 1} \frac{\tilde{u}_i(s^*)}{1 - p_i(s^*)}. \end{aligned}$$

Thus, $Y_j(s^*) = 0$ implies $\tilde{u}_i(s^*) = 0$. 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 (2) that $u^* > 0$. □

PROOF OF PROPOSITION 10. (i) *No dominant signals implies the existence of interior limit.* Define

$$L^0 = \{s \in [0, 1]^n \mid \tilde{u}_i(s) \leq 0 \text{ for all } i, \text{ with equality for at least one } i \in N\}.$$

By aggregate uncertainty, $L^0 \neq \emptyset$. For any $s \in L^0$, set $I(s) = \{i \in N \mid \tilde{u}_i(s) = 0\}$. If $|I(s)| = 2$ and $s \in (0, 1)^n$, then s is an interior limit. Suppose there is no interior limit, so that $|I(s)| = 1$ for all $s \in L^0 \cap (0, 1)^n$. By continuity of payoffs in signals, this means that there exists $i \in N$ such that $I(s) = i$ for all $s \in L^0 \cap (0, 1)^n$. Hence $\tilde{u}_i(s) > \tilde{u}_j(s)$ for all $s \in L^0 \cap (0, 1)^n$.

Consider first the case in which all agents are strongly optimistic. Take $s = (s_1, \dots, s_n) \in L^0$ with $s_i = 0$ and $s_j > 0$ for all $j \neq i$ (which exists by aggregate uncertainty). By strong optimism and MLRP, we have $h_i(s_i) = 1$ and $h_j(s_j) > 1$, and thus

$$\tilde{u}_i(s) = \alpha(s)h_i(s_i)H - 1 < \alpha(s)h_j(s_j)H - 1 = \tilde{u}_j(s),$$

which contradicts the finding that $\tilde{u}_i(s) > \tilde{u}_j(s)$ on $L^0 \cap (0, 1)^n$ (using the continuity of \tilde{u}_i).

Second, consider the case in which some agent i is not strongly optimistic. Then there exists a signal profile $s = (s_1, \dots, s_n)$ with $s_i > 0$ and $s_j = 1$ for all $j \neq i$, such that $\tilde{u}_i(s) = 0$. Since $I(s) = i$ for all $s \in L^0 \cap (0, 1)^n$, we have $\tilde{u}_i(s') > \tilde{u}_j(s')$ for all $s' \in L^0 \cap (0, 1)^n$, so that by continuity of \tilde{u} , $0 = \tilde{u}_i(s) \geq \tilde{u}_j(s)$. But then s is a dominant signal for agent i , contradicting the hypothesis.

(ii) Suppose player i is strongly optimistic. By aggregate uncertainty, there exists a nonempty set of signal profiles

$$L_i = \{s \in [0, 1]^n \mid s_i = 0, \tilde{u}_i(s) = 0\}.$$

Let $u^* = \max_{j \neq i, s \in L_i} \tilde{u}_j(s)$ (which exists because L_i is compact and \tilde{u}_j is continuous for all j). Denote the solution by u^* and the maximizing arguments by j^* and s^* . Then, by

construction, $s_i^* = 0$, $\tilde{u}_i(s^*) = 0$, $\tilde{u}_{j^*}(s^*) = u^*$, and $\tilde{u}_j(s^*) \leq u^*$ for all $j \neq i, j^*$. It is easy to verify that $Y_i(s^*) = Y_{j^*}(s^*) = 0$. Thus, s^* is a boundary limit. \square

REFERENCES

- Aidt, Toke S. and Peter S. Jensen (2009), “The taxman tools up: An event history study of the introduction of the personal income tax.” *Journal of Public Economics*, 93, 160–175. [342]
- Banerjee, Abhijit V. (1992), “A simple model of herd behavior.” *Quarterly Journal of Economics*, 107, 797–817. [343]
- Becker, Marshall H. (1970), “Factors affecting diffusion of innovations among health professionals.” *American Journal of Public Health and the Nations Health*, 60, 295–304. [342]
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), “A theory of fads, fashion, custom, and cultural change as informational cascades.” *Journal of Political Economy*, 100, 992–1026. [343, 361, 363]
- Bliss, Christopher and Barry Nalebuff (1984), “Dragon-slaying and ballroom dancing: The private supply of a public good.” *Journal of Public Economics*, 25, 1–12. [343, 344]
- Bolton, Patrick and Christopher Harris (1999), “Strategic experimentation.” *Econometrica*, 67, 349–374. [344]
- Chamley, Christophe and Douglas Gale (1994), “Information revelation and strategic delay in a model of investment.” *Econometrica*, 62, 1065–1085. [343, 344, 363]
- Daron, Acemoglu, Kostas Bimpikis, and Asuman Ozdaglar (2011), “Experimentation, patents, and innovation.” *American Economic Journal: Microeconomics*, 3, 37–77. [344]
- Décamps, Jean-Paul and Thomas Mariotti (2004), “Investment timing and learning externalities.” *Journal of Economic Theory*, 118, 80–102. [344]
- Eeckhoudt, Louis and Christian Gollier (1995), “Demand for risky assets and the monotone probability ratio order.” *Journal of Risk and Uncertainty*, 11, 113–122. [347]
- Fudenberg, Drew, David M. Kreps, and David K. Levine (1988), “On the robustness of equilibrium refinements.” *Journal of Economic Theory*, 44, 354–380. [353]
- Fudenberg, Drew and Jean Tirole (1991), *Game Theory*. MIT Press, Cambridge, Massachusetts. [351]
- Goolsbee, Austan and Peter J. Klenow (2002), “Evidence on learning and network externalities in the diffusion of home computers.” *The Journal of Law and Economics*, 45, 317–343. [342]
- Gul, Faruk and Russell Lundholm (1995), “Endogenous timing and the clustering of agents’ decisions.” *Journal of Political Economy*, 103, 1039–1066. [344]
- Herrera, Helios and Johannes Hörner (2013), “Biased social learning.” *Games and Economic Behavior*, 80, 131–146. [352]

Hofbauer, Josef and Karl Sigmund (1998), *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge. [364]

Kaustia, Markku and Samuli Knüpfer (2012), “Peer performance and stock market entry.” *Journal of Financial Economics*, 104, 321–338. [342]

Keller, Godfrey, Sven Rady, and Martin Cripps (2005), “Strategic experimentation with exponential bandits.” *Econometrica*, 73, 39–68. [344]

Liu, Hong, Qi Sun, and Zhong Zhao (2014), “Social learning and health insurance enrollment: Evidence from China’s new cooperative medical scheme.” *Journal of Economic Behavior and Organization*, 97, 84–102. [342]

Malueg, David A. and Shunichi O. Tsutsui (1997), “Dynamic R&D competition with learning.” *The RAND Journal of Economics*, 28, 751–772. [344]

Moscarini, Giuseppe and Francesco Squintani (2010), “Competitive experimentation with private information: The survivor’s curse.” *Journal of Economic Theory*, 145, 639–660. [344]

Murto, Pauli and Juuso Välimäki (2011), “Learning and information aggregation in an exit game.” *The Review of Economic Studies*, 78, 1426–1461. [344, 364]

Murto, Pauli and Juuso Välimäki (2013), “Delay and information aggregation in stopping games with private information.” *Journal of Economic Theory*, 148, 2404–2435. [344]

Olson, Mancur (1965), *The Logic of Collective Action*. Harvard University Press, Cambridge, Massachusetts. [363]

Palfrey, Thomas R. and Howard Rosenthal (1984), “Participation and the provision of discrete public goods: A strategic analysis.” *Journal of Public Economics*, 24, 171–193. [341, 343]

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2007), “Social learning in one-arm bandit problems.” *Econometrica*, 75, 1591–1611. [344, 364]

Co-editor Nicola Persico handled this manuscript.

Manuscript received 5 May, 2015; final version accepted 27 October, 2016; available online 10 February, 2017.