GOVERNANCE AND THE EFFICIENCY
OF ECONOMIC SYSTEMS
GESY

Discussion Paper No. 500

Who goes first? Strategic Delay and Learning by Waiting

Peter Wagner *

* University of Bonn

May 2015

Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 is gratefully acknowledged.

Sonderforschungsbereich/Transregio 15 · www.sfbtr15.de
Universität Mannheim · Freie Universität Berlin · Humboldt-Universität zu Berlin · Ludwig-Maximilians-Universität München
Rheinische Friedrich-Wilhelms-Universität Bonn · Zentrum für Europäische Wirtschaftsforschung Mannheim

Speaker: Prof. Dr. Klaus M. Schmidt · Department of Economics · University of Munich · D-80539 Munich,
Phone: +49(89)2180 2250 · Fax: +49(89)2180 3510

Who goes first? Strategic Delay and Learning by Waiting*

Peter Wagner†

May 4, 2015

Abstract

This paper considers a "war of attrition" game in which agents learn about an uncertain state of the world through private signals and from their peers. I provide existence and uniqueness results for a class of equilibria that satisfy a "full-participation" condition, and show that asymmetries in the distribution of information can lead to excessive stopping and an oversupply of information relative to the social optimum.

1 Introduction

In 1929, the young German physician Werner Forssmann secretely conducted a risky self-experiment. He inserted a narrow tube into his arm and maneuvered it along a vein unto his heart. The procedure, known as cardiac catheterization, constituted a revolutionary breakthrough in cardiology, and

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

†Department of Economics, University of Bonn, Germany

later earned him the Nobel prize in medicine. Forssmann's main contribution was the proof that cardiac catheterization was safe to perform on humans. The basic methods for the procedure had already been developed decades earlier and successfully tested on animals. It was widely believed, however, that inserting any object into the beating human heart was fatal, and thus there was a need for someone to put this hypothesis to the ultimate test.

The story of Werner Forssman is of someone who took action in an environment of "wait and see", in which everyone hoped for the independent initiative of a volunteer who resolves some of the risks relating to an uncharted course of action. The fact that people learn from the behavior and experience of their peers has been shown empirically by a large body of literature. Peer learning effects have been found for example in the diffusion of innovations among health professionals (Becker, 1970), the enrollment in health insurance (Liu et al., 2014), the diffusion of home computers (Goolsbee et al., 2002), stock market entry (Kaustia and Knüpfer, 2012) and the introduction of the personal income tax (Aidt and Jensen, 2009). In environments in which no formal institution or informal arrangement exists that coordinates exploratory activities, how efficient is it to rely on the initiative of volunteers, and how well does such a decentralized mechanism aggregate dispersed information?

To study this problem, we consider a stopping game with asymmetric distribution of information and a pure informational externality. In this game, each agent has the option to take the same risky action. Taking this action produces a payoff which depends on an uncertain state of the world. At the beginning of the game agents privately receive information about the state and then they engage in a "war of attrition", each deciding independently how long to wait before taking action. The first agent who stops receives his state-dependent payoff and thereby reveals the state to the remaining agents. Uncertainty about the value of taking action and payoff observability generate a second-mover advantage that provides agents with an incentive to free-ride on others' initiative.

I allow for heterogeneity in the distribution of information among agents. Existing literature dealing with social learning and free-riding typically considers symmetric equilibria of models in which agents are ex-ante identical. Because of the heterogeneity in my model, we cannot use symmetric equilibrium. Instead, I propose a different solution concept, called *full-participation equilibrium*, in which the (ex-ante) distribution over stopping times induced by each agent's strategy has full support. Intuitively, the full-participation requirement rules out equilibria in which agents use the threat of being passive in the future to coerce others into taking action. I prove uniqueness of the full-participation equilibrium and its existence for a general class of games. Moreover, I demonstrate that the full-participation equilibrium coincides with the symmetric equilibrium when agents are identical. As such, full-participation equilibrium can be viewed as a natural generalization of symmetric equilibrium for models with heterogeneous agents.

I analyze the effects of variations in the distribution of information on the full-participation equilibrium outcome and show that a strong informational asymmetry leads to excessive stopping by poorly informed agents. More specifically, when there are two agents of whom one is significantly better informed than the other, then in a full-participation equilibrium the agent with the least accurate information stops with probability one. Therefore, the state is revealed with certainty regardless of the realization of signals, which is more than would be optimal from a social welfare point of view. The cost is borne mainly by the poorly informed.

There is a number of results in the related literature that suggest the opposite: when a public good is provided through voluntary contribution, then in a symmetric equilibrium the good is provided at an socially insufficient level, because no agent takes into account the positive externality of his own contribution for others. Our model differs from these standard public goods papers in two important ways. First, agents have asymmetrically distributed private information and learn from others' inaction about the state over time. Second, there is aggregate uncertainty whether providing the good is individually and socially desirable. Naively, one may expect that the standard logic still applies by the argument that when one agents observes that others remain passive, she becomes more pessimistic about the state and is thus less willing to take action herself. In fact, this is indeed the case when information is symmetrically distributed.

The argument does not go through however when there is a strong asymmetry in the distribution of information. The basic intuition is that less informed agents crowd out the informed agents' participation.To fix ideas, consider a model with two groups of agents. The first group observes informative signals about the state from some common distribution, the second group receives no information. In equilibrium each agent chooses a stopping time based on his belief about others' behavior. The informed agents' willingness to take action correlates with the state of the world: they tend to be more pessimistic when the state is low and are thus tend to wait longer. The behavior of the uninformed, on the other hand, is uncorrelated with the state of the world. This means that the informed have an informational reason to wait when they receive a low signal, as well as a strategic reason to do so, resulting from the relatively higher optimism of the uninformed. In equilibrium, the informed stop at decreasing rates over time, and in return, the uninformed learn less and less from the informed agents' inactivity. In the limit, as time goes to infinity, informed agents become entirely inactive and thus their abstinent behavior provides no information to the remaining agents.

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

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

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

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

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

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

2 Model

The set of agents is $N = \{1, \ldots, n\}$. Each agent decides when to take a risky irreversible action that yields the uncertain payoff that depends on an uncertain state of the world $\theta \in \{H, L\}$, where $H > 0$ is arbitrary and L is normalized to -1. Once one agent stops, the remaining agents can observe the true value of θ and they can choose to collect the payoff θ as well or to make use of an outside option that yields a payoff of zero. There is no cost to staying in the game. Agents share the common prior belief p_0 that $\theta = H$. At the outset, it is commonly known that each agent $i \in N$ is endowed with a signal $s_i \in S_i = [\underline{s}_i, \bar{s}_i]$ about θ, drawn independently from a distribution which has full support and a smooth c.d.f. $F_{i,\theta}$ with bounded density. Future payoffs are discounted at common rate r.

Strategies and equilibrium. A pure strategy for agent $i \in N$ is a function $\sigma_i : [\underline{s}_i, \bar{s}_i] \to [0, \infty) \cup \{\infty\}$, where $\sigma_i(s_i)$ represents the time agent i with signal s_i stops if the game has not ended. For a given profile of signals and strategies, denote by t_{-i} the first taking action time among all agents except i. If agent i chooses $t_i > t_{-i}$, then agent i observes the state of the world at time t_{-i}, and consequently receives the payoff $e^{rt_{-i}} \max\{\theta, 0\}$. If agent i stops at a time $t_i \leq t_{-i}$, she receives a payoff $e^{rt_{-i}}\theta$. For a given profile $\sigma_{-i} = (\sigma_j)_{j \neq i}$ of strategies for all agents except i, the expected payoff for agent i with signal s_i from taking action at time t_i is

$$U_i(t_i, \sigma_{-i}|s_i) = \Pr(t_{-i} < t_i|s_i)\mathbb{E}[e^{-rt_{-i}} \max\{\theta, 0\}|t_{-i} < t_i, s_i] \\ + \Pr(t_{-i} \geq t|s_i)\mathbb{E}[e^{-rt_i}\theta|t_{-i} \geq t_i, s_i].$$

The expectations are with respect to θ and t_{-i}. A strategy profile $(\sigma_i)_{i \in N}$ is a Bayes-Nash-equilibrium if $\sigma_i(s_i) \in \arg\max_{t_i} U_i(t_i, \sigma_{-i}|s_i)$ for all i and every $s_i \in [\underline{s}_i, \bar{s}_i]$.

We consider *full-participation equilibria* which are Bayes-Nash equilibria in which the image of each strategy is the entire positive real line.

Definition 1 (Full-participation equilibrium). *A pure strategy Bayes-Nash equilibrium $(\sigma_i)_{i \in N}$ is a full-participation equilibrium if for all $t \in \mathbb{R}_+$ and every $i \in N$ there exists $s_i \in S_i$ such that $\sigma_i(s_i) = t$.*

We prove that full-participation equilibria have a number of desirable properties. They are unique, and exist in a large class of games. Moreover, strategies that are part of a full-participation equilibrium are differentiable and monotone (see Lemma 1 on page 12). More specifically, if σ_i is a strategy belonging to a full-participation equilibrium, then there is a $\hat{s}_i \in (\underline{s}_i, \bar{s}_i)$ such that σ_i is constant on $[\underline{s}_i, \hat{s}_i]$, and strictly decreasing on $(\hat{s}_i, \bar{s}_i]$. Moreover, when agents are identical, then symmetric equilibrium and full-participation equilibrium coincide. In this sense, the notion of full-participation equilibrium is a natural analogue to symmetric equilibrium when there is heterogeneity among agents.

Note that assuming that agents choose their taking action time simultaneously ex-ante is without loss of generality. We could write the model as a full-fledged dynamic game, in which agents decide continually over time whether to take action or to wait and continuously update their beliefs about the state of the world and the other agents' signals. By construction, every history in a full-participation equilibrium is reached along the equilibrium path, and thus each full-participation equilibrium of the dynamic game is outcome-equivalent to the unique full-participation equilibrium of the corresponding game in strategic form.

Signals. We assume that signal distributions have the monotone likelihood ratio property (MLRP), that is, the likelihood ratio $F'_{i,H}(s_i)/F'_{i,L}(s_i)$ is increasing in s_i for each agent i. We shall make two further assumptions to render the strategic interaction interesting. The first assumption ensures each agent is willing to take action in equilibrium for some realizations of his signal.

Assumption 1 (Initial Optimism)**.** $\mathbb{E}[\theta|\bar{s}_i] > 0 \quad \forall i \in N.$

Initial optimism among all agents is a necessary condition for the existence of full-participation equilibria, since agents for whom the assumption does not hold would never act in any equilibrium. We can transform any game into one that satisfies initial optimism by removing all agents for whom the expected value of the state for each of his signals is negative.

Next, we assume that there is aggregate uncertainty about the state of the world. By aggregate uncertainty we mean that there is a profile of signals such that the expected value of the state conditional on this signal profile is negative. It follows that there is a signal for each agent so that this agent prefers not to act for some realization of other agents' signals.

Assumption 2 (Aggregate Uncertainty)**.**

$$\Pr(H|\underline{s}_1, \ldots, \underline{s}_n)nH - \Pr(L|\underline{s}_1, \ldots, \underline{s}_n) < 0.$$

Aggregate uncertainty in conjunction with initial optimism implies that it is initially uncertain whether taking action is socially desirable. The assumption of aggregate uncertainty is important, because we are interested in studying issues relating the aggregation of dispersed information, and want to assess how the equilibrium outcome in the strategic environment compares to the socially optimal outcome. Without aggregate uncertainty, taking action would always be socially optimal, so that questions relating to efficiency at the extensive margin become moot.

The *reverse hazard rate* $\lambda_{i,\theta}$ of agent i's at $s_i \in S_i$ in state θ is the density assigned to signal s_i, conditional on the event that the true signal is weakly lower than s_i. It is given by

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

The *reverse hazard rate ratio* h_i for agent i at $s_i \in S_i$ is defined as the ratio of reverse hazard rates, and given by

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

It is well known that MLRP implies $\lambda_{i,H} > \lambda_{i,L}$ and thus $h_i > 1$.

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

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

MLRP implies that the $F_{i,h}/F_{i,L}$ is increasing for each i (Eeckhoudt and Gollier, 1995) and thus α is increasing in each of its arguments.

The hazard rate ratio h_i and the likelihood ratio of the public posterior α play an important role in the equilibrium analysis. If $(\sigma_i)_{i\in N}$ is a full-participation equilibrium, then each σ_i is monotone (see Lemma 1 on page 12), and thus there exists an inverse ϕ_i of σ_i restricted to some subinterval $(\hat{s}_i, \bar{s}_i]$ of S_i. Therefore, at time $t = \sigma_i(s_i)$, the posterior probability agent i with signal $s_i \in (\hat{s}_i, \bar{s}_i]$ assigns to the state being H has the likelihood ratio

$$\frac{p_0}{1-p_0} \prod_{j\neq i} \frac{F_{j,H}(\phi_j(t))}{F_{j,L}(\phi_j(t))} \frac{F'_{i,H}(\phi_i(t))}{F'_{i,L}(\phi_i(t))} = \alpha(\phi_1(t), \ldots, \phi_n(t)) h_i(\phi_i(t)).$$

This equation says that in an equilibrium in monotone strategies, for each time t the likelihood ratio of the subjective belief of an agent i who is to act at time t can be decomposed into the common component α and a private component h_i. Here, α represents a measure of the information about the state that is commonly available to all agents. The factor h_i represents the information that agent i holds privately and it provides a measure of divergence of an agent's private belief from the public belief. Based on these definitions, we introduce the following notion of asymmetry in the distribution of information between agents.

Definition 2. *The distribution of information is strongly asymmetric if there are $i, j \in N$ such that* $\inf_{s_i\in S_i} h_i(s_i) > \sup_{s_j\in S_j} h_j(s_j)$.

In other words, the distribution of signals is strongly asymmetric, if there are agents i and j so that the lowest value of h_i exceeds the highest value of h_j. The more informative a signal is, the larger the divergence of an agents subjective belief about the state from the public belief. Consequently, higher values of the reverse hazard rate ratios are associated with more informative signals. Thus, strong asymmetry in the distribution of information essentially means that there is a pair of agents such that one is always better informed than the other, regardless of the actual realization of signals and the nature of the equilibrium strategies.

As an example, consider a signal distribution that has the property $F_{i,H} = F_{i,L}^{\beta_i}$, where $\beta_i > 1$. Such a signal distribution has a constant reverse hazard rate ratio equal to β_i. Moreover, the informativeness of signals is strictly increasing in β_i. Thus, if there are two agents i and $j \neq i$, whose signal distributions have the constant reverse hazard rate ratios β_i and $\beta_j < \beta_i$, then the distribution of information is strongly asymmetric, with agent i being better informed than agent j.

3 Equilibrium existence and uniqueness

First, we derive conditions for a strategy profile to be a full-participation equilibrium and we then we how that these condition imply uniqueness. Second, we provide sufficient conditions for existence.

3.1 Uniqueness

Fix an agent i and a strategy for every other agent $j \neq i$. Denote by $G_{i,\theta}(t)$ the probability that agent i assigns to the event that every other agent chooses an action lower than t, conditional on the state being θ. 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 given by

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

Then, agent i's expected payoff is

$$(1) \quad u_i(t, s_i) = \Pr(H|s_i) \left(\int_0^t e^{-rz} dG_{i,H}(z) + (1 - G_{i,H}(t)) e^{-rt} \right) H \\ - \Pr(L|s_i)(1 - G_{i,L}(t)) \, e^{-rt}.$$

The first term in (1) is the expected payoff from taking action at t conditional on the state being high. Agent i with signal s_i assigns probability $\Pr(H|s_i)$ to this event. He receives payoff $e^{-rz}H$ if another agent acts at $z < t$, and otherwise he acts at time t and obtains the payoff $e^{-rt}H$. The second term represents the expected payoff if the state is low. In this case agent i receives a payoff of zero if some agent acts before t, and otherwise he incurs a loss $-e^{-rt}$. We first show that every full-participation equilibrium is in monotone and differentiable strategies.

Lemma 1. *Let $(\sigma_1, \ldots, \sigma_n)$ be a full-participation equilibrium. Then for each σ_i there exist $\hat{s}_i \in [\underline{s}_i, \bar{s}_i]$, so that*

(i) $\sigma_i(s_i) = \infty$ *for all* $s_i \in [\underline{s}_i, \hat{s}_i]$.

(ii) σ_i *is strictly decreasing and differentiable on* $(\hat{s}_i, \bar{s}_i)$.

(iii) $\sigma_i(\bar{s}_i) = 0$.

The proof is in the appendix. The lemma allows us to write each $G_{i,\theta}$ as a function of the signal distributions of all agents $j \neq i$ and their strategies. Specifically, there exists an inverse ϕ_i to each σ_i on $(\hat{s}_i, \bar{s}_i]$ which itself is strictly decreasing. Therefore, the probability that the earliest stopping time among all agents except i is after time t is equal to the joint probability that the signal of each agent $j \neq i$ is below $\phi_j(t)$, so that by conditional independence

$$1 - G_{i,\theta}(t) = \prod_{j \neq i} F_{j,\theta}(\phi_j(t)). \tag{2}$$

By part (ii) of the lemma and smoothness of the distribution functions, each $G_{i,\theta}$ is differentiable. It thus follows immediately from (2) that

$$\frac{G'_{i,\theta}(t)}{1 - G_{i,\theta}(t)} = -\left(\sum_{j \neq i} \frac{F'_{j,\theta}\left(\phi_j(t)\right)}{F_{j,\theta}\left(\phi_j(t)\right)} \phi'_j(t)\right). \tag{3}$$

Taking the partial derivative of equation (1) with respect to t and setting it equal to zero yields

$$\Pr(H|s_i)(1 - G_{i,H}(t))rH - \Pr(L|s_i)\left(G'_{i,L}(t) + r(1 - G_{i,L}(t))\right) = 0. \tag{4}$$

Now, divide both sides by $r \Pr(L|s_i)(1 - G_{i,L}(t))$ and rewrite the last expression as

$$\frac{\Pr(H|s_i)}{\Pr(L|s_i)} \frac{1 - G_{i,H}(t)}{1 - G_{i,L}(t)} H - 1 = \frac{1}{r} \frac{G'_{i,L}(t)}{1 - G_{i,L}(t)}. \tag{5}$$

If we now substitute (2) and (3) in (5), we obtain

$$\frac{p_0}{1 - p_0} \frac{F'_{i,H}(\phi_i(t))}{F'_{i,L}(\phi_i(t))} \prod_{j \neq i} \frac{F_{j,H}(\phi_j(t))}{F_{j,L}(\phi_j(t))} H - 1 = -\frac{1}{r}\left(\sum_{j \neq i} \frac{F'_{j,L}(\phi_j(t))}{F_{j,L}(\phi_j(t))} \phi'_j(t)\right).$$

Finally, we can write the last equation in terms of the revere hazard rates $\lambda_{i,\theta}$, the reverse hazard rate ratios h_i and the likelihood ratio of the posterior α, and obtain the more succinct expression

$$h_i(\phi_i(t))\alpha(\phi(t))H - 1 = -\frac{1}{r} \sum_{j \neq i} \lambda_{j,L}(\phi_j(t))\phi'_j(t), \tag{6}$$

where $\phi = (\phi_1, \ldots, \phi_n)$ denotes the profile of inverse strategies. Using elementary operations to isolate the derivatives in the equations for $i = 1, \ldots, n$, we obtain the nonlinear differential equation

$$(7) \quad -\phi_i'(t) = \frac{r}{\lambda_{i,L}(\phi_i(t))}\left(\alpha(\phi(t))\left[\frac{1}{n-1}\sum_{j=1}^{n} h_j(\phi_j(t)) - h_i(\phi_i(t))\right]H - \frac{1}{n-1}\right).$$

for each $i = 1, \ldots, n$. In a full-participation equilibrium, the highest type of each agent acts immediately (by Lemma 1), so that a solution of (7) corresponds to a full-participation equilibrium if it is strictly decreasing and $\phi(0) = (\bar{s}_1, \ldots, \bar{s}_n)$. The next lemma shows that the dynamic system (7) has a unique solution.

Lemma 2. *The system of ordinary differential equations* (7) *with initial condition* $\phi(0) = (\bar{s}_1, \ldots, \bar{s}_n)$ *has a unique solution.*

The proof is standard, and based on the fundamental existence and uniqueness theorem for Lipschitz-continuous systems of ordinary differential equations. The lemma implies that if a full-participation equilibrium exists, then it is unique. The solution to (7) is not necessarily strictly monotone, however, so that a full-participation equilibrium may fail to exist in general.

3.2 Sufficient conditions for existence

In this subsection we provide sufficient conditions for the existence of full-participation equilibria. We consider games with $n \geq 2$ agents of whom $m \in N$ receive a signal drawn from a distribution F_θ in state $\theta \in \{H, L\}$, with reverse hazard rate ratio h and reverse hazard rate λ_θ (omitting the index i). Except for differentiability and full support we do not impose restrictions on F_θ. The remaining $n - m$ agents receive state-independent signal. For expositional purposes, I refer to the agents who receive informative signals as "experts", and I call those agents who receive no information "amateurs".

We write

$$\alpha_m(s) = \frac{p_0}{1-p_0}\left(\frac{F_H(s)}{F_L(s)}\right)^m .$$

This is the likelihood ratio of the experts' posterior belief about θ when each of the experts' signal is equal to s.

The equilibrium strategy for experts is obtained from (7), using the fact that the differential equation for amateurs depends only on the experts' strategy. This specification allows us to derive the strategy for experts in closed form. The strategy for amateurs on the other hand is not available in closed form, but we can characterize the amateurs' *stopping rate* which is the hazard rate of the distribution of their individual stopping times. More specifically, if ϕ_0 denotes the strategy of an amateur and F_0 the state independent distribution of the amateurs signal, then the stopping time has the c.d.f. $G_0(t) = 1 - F_0(\phi_0(t))$. The associated stopping rate is therefore given by $\lambda_0(\phi_0(t))\phi_0'(t)$. The following proposition shows that a full-participation equilibrium exists.

Proposition 1. *Let $n \geq 2$ with $n \geq m \geq 0$ informed and $n-m$ uninformed agents. There exists a full-participation equilibrium if*

$$p_0[n-m-(n-m-1)h(\bar{s}))]H > 1 - p_0 \tag{8}$$

in which the informed agents' stopping time is

$$\sigma(s) = \int_s^{\bar{s}} \frac{(n-1)\lambda_L(z)/r}{\alpha_m(z)\left[n-m-(n-m-1)h(z)\right]H-1} dz$$

for an agent with signal $s > s^*$ and $\sigma(s) = 0$ if $s \leq s^*$, where s^* solves

$$\alpha_m(s)(n-m-(n-m-1)h(s))H = 1. \tag{9}$$

The strategy of each uninformed agent induces the stopping rate

$$\frac{r}{n-1}\Big(\alpha_m(\phi(t))\big[(h(\phi(t))-1)m+1\big]H-1\Big) \tag{10}$$

where ϕ is the inverse of σ.

The above equilibrium specification includes as special cases the symmetric games in which (i) agents receive identically distributed signals ($m = n$) and (ii) no agent has any information ($m = 0$). In both cases, the equilibrium is a symmetric one. Note that for $m = n$, condition (8) is satisfied because of Assumption 1. For $m = 0$, when all signals are drawn from a state-independent distribution, then (8) reduces to $p_0 H > 1 - p_0$ which simply means that the stopping value must be positive for each agent. The uninformed agents' stopping rate then becomes

$$\frac{r}{n-1}\left(\frac{p_0}{1-p_0}H-1\right).$$

Assumption 2 is violated in that case. When $p_0\,H < 1 - p_0$, then in the unique equilibrium no agent stops in finite time, and thus no full-participation equilibrium exists. In all other cases ($n > m > 0$), we have existence of a full-participation equilibrium only if the arrival rate at the highest signal is low enough, which roughly means that the best signal informed agents receive cannot be too informative. The bound depends on the number of agents n, on of the number of informed agents m, and on the relative gain H when the state is high. Condition (8) is needed, because full-participation equilibria may fail to exist if the signal structures are extremely skewed, with some agents having significantly more informative signals than others. It is then possible that no strategy profile exists such that all agents are willing to participate.

As an illustration, consider a situation with one expert (agent 1, say) and two amateurs, so that $h_1 > 1$, and $h_2 = h_2 = 1$. In this case, we have $\frac{1}{2}\sum_{j=1}^{3} h_j(s_j) - h_1(s_i) = 1 - h_1(s_1)/2$, which is negative if $h_1(s_1) > 2$. Thus, if agent 1's signal is very informative, from equation (7), it follows that $\phi_1'(\epsilon) > 0$ for $\epsilon > 0$ small, ruling out that agent $1's$ strategy is strictly decreasing. The underlying intuition is as follows. If agent i receives a signal $s_i = \bar{s}_i - \epsilon$, he must choose an action lower than 1 by strict monotonicity of strategies by Lemma 1. From equation (6) it follows that inducing agent 1 with signal s_1 to concede at $t > 0$ requires that the stopping rates for agents 2 and 3 are sufficiently high. But both agents 2 and 3 are less optimistic about the state than agent 1, so that if agent 2's strategy makes it a best response for agent 1 to choose a lower action, then agent 3 will strictly prefer a later stopping time regardless of his type, implying that agent 3's strategy must have a jump, thus ruling out existence of a full-participation equilibrium.

4 Efficiency

We now introduce a notion of efficiency that will serve as a benchmark of comparison for the equilibrium analysis. In principle, there are two types of inefficiencies that can arise in equilibrium. First, inefficiencies can arise at the intensive margin, because agents act only with delay as a result of free-riding. This issue has been discussed in detail in Bliss and Nalebuff (1984). We thus set aside this aspect and focus on the inefficiencies along the extensive margin, resulting from too much or too little stopping as a consequence of imperfect exchange of private information.

Our notion of efficiency focuses on the aggregate stopping decision and entails the constraint that each agent's strategy cannot depend on another agent's strategy. Intuitively, this corresponds to a situation in which agents collaboratively choose a profile of strategies *prior* to observing their signals so as to maximize the sum of their payoffs. Let the *stopping region* for a given strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a subset $E \subset S_1 \times \ldots \times S_n$ of signal profiles defined as

$$E = \left\{ (s_1, \ldots, s_n) \middle| \min_i \{\sigma_i(s_i)\} < \infty \right\}.$$

The stopping region is the set of all signal profiles for which some agent takes action in finite time. We shall say that for a given strategy profile σ aggregate stopping is *efficient* if the stopping region E associated with σ maximizes the sum of payoffs resulting from one agent taking action immediately if and only if the realized signal profile lies in E. We call such a stopping region the efficient stopping region.

The efficient stopping region is characterized by thresholds, one threshold $\hat{s}_i$ for each agent i. A signal profile $(s_1, \ldots, s_n)$ then lies in the stopping region if and only if for any i, agent i's signal s_i lies above the threshold $\hat{s}_i$. This follows from the monotone likelihood ratio property: if it is socially optimal for an agent take action in finite time when his signal is s_i, then it must also be socially optimal to take action for any signal $s_i' > s_i$ as the higher signal implies a higher probability of a high state. The socially optimal profile of thresholds is chosen such that social gain from taking action is maximized, conditional on agent i's signal and on the event that i is pivotal. More specifically, the efficient stopping region is characterized as follows.

Proposition 2. *The efficient stopping region is given by the set of signal profiles* $\{(s_1, \ldots, s_n) \mid \exists\, i \in N : s_i \geq \hat{s}_i\}$*, where* $\hat{s}_i > \underline{s}_i$ *for each* $i \in N$ *and each* $\hat{s}_i < \bar{s}_i$ *solves*

$$\alpha(\hat{s}_1, \ldots, \hat{s}_n) h_i(\hat{s}_i) = \frac{1}{nH}.$$

Note that the efficient stopping region differs from the set of signal profiles at which a social planner would stop who can observe the agents' signals. A benchmark involving such a omniscient social planner would set an excessively high standard for efficiency as it demands that agents use strategies that are measurable with respect to all available information. The stopping region that is optimal for a social planner is thus never attainable in equilibrium.

It is interesting to note that when there is an agent who is more informed than all others in the sense of Definition 2, then this agent's signal alone determines whether action is taken or not. More specifically, suppose we have n agents, and suppose the lowest value of the reverse hazard ratio of agent 1 is higher than the highest value of the reverse hazard rate ratio of each agent $i > 1$. Proposition 2 then implies that the signal profile $(s_1, \ldots, s_n)$ lies in the efficient stopping region if and only if $s_1 > \hat{s}_1$, where $\hat{s}_1$ is the signal of agent 1 for which the expected social value of taking action is equal to zero. The fact that the most informed agent becomes pivotal follows immediately from the condition in Proposition 2. When there is a gap in the reverse hazard ratios, then the condition cannot hold simultaneously for more than one agent.

5 Asymmetric distribution of information and excessive stopping

In this section, we analyze the effects of variations in the distribution of information on the full-participation equilibrium outcome and show that strong informational asymmetry leads to excessive stopping by poorly informed agents. More specifically, when there are two agents of whom one is significantly better informed than the other, then in a full-participation equilibrium the agent with the least accurate information stops with probability one. Therefore, the state is revealed with certainty regardless of the realization of signals

Theorem 1. *If the distribution of information is strongly asymmetric, then in the full-participation equilibrium there exists an agent who stops in finite time almost surely.*

The theorem tells us that information aggregation completely fails when there are asymmetries in the distribution of information. Moreover, the inefficiency does not arise through insufficient stopping due to free-riding as has been suspected in the literature (Foster and Rosenzweig, 2010), but rather from *excessive* stopping. The cost of this excess is borne mainly by the poorly informed. Thus, as far as production of information is concerned, there is no "exploitation of the great by the small" (Olsen, 1965) but rather an exploitation of the small by the great.

To illustrate the logic behind the result, let us consider the game from the last subsection with m experts and $n-m$ amateurs where $n > m > 0$. First note that the distribution of information is strongly asymmetric, since MLRP implies that the reverse hazard rate ratio for experts is greater than one, while it is equal to one for amateurs. From (9) it follows immediately that the right-hand side of (7) converges to zero as t approaches infinity. Hence, the stopping rate of experts converges to zero. Moreover, by Assumption 2, the threshold s^* defined in (9) is greater than the experts' lowest signal, which implies that they stop with probability less than one.

Next, consider the stopping rate for amateurs. By Lemma 1, equilibrium strategies are strictly decreasing, and thus $\alpha_m(\phi(t))$ is strictly decreasing in t. Moreover, the reverse hazard rate ratio for experts' signal distribution is strictly greater than 1, and thus its infimum, denoted by h^*, is greater than 1 as well. Hence, letting $\alpha_m^* = \lim_{t\to\infty} \alpha_m(\phi(t))$, the stopping rate of amateurs at any $t \in [0, \infty)$ is larger than

$$\frac{r}{n-1}\Big(\alpha_m^*\big((h^*-1)m+1\big)H - 1\Big). \tag{11}$$

From the equilibrium threshold defined in (9), it follows that

$$1 \leq \alpha_m^*(n-m-(n-m-1)h^*)H.$$

If we now substitute this inequality into (11), we find that the stopping rate for amateurs has the lower bound

$$\alpha_m^*(h^* - 1)H > 0.$$

This last inequality implies that the stopping rate of amateurs is greater than zero for all $t > 0$. Moreover, it shows that the distribution over stopping times induced by the amateurs' strategy first-order stochastically dominates the exponential distribution with parameter $\alpha_m^*(h^* - 1)H$. Consequently, amateurs stop in finite time with probability one.

In summary, in a full-participation equilibrium of a game with informed "experts" and uninformed "amateurs", the experts' stopping rates converges to zero as t tends to infinity, while the uninformed agents continue to stop at a positive rate throughout. It is noteworthy that amateurs as well as some informed experts retain a positive stopping value. That is, their expected instantaneous payoff from stopping remains positive even in the limit as t approaches infinity. This is clear for amateurs, as they would otherwise not be willing to stop. For experts, note that an expert whose signal is at the equilibrium threshold s^* is more optimistic about the state than an amateur in the limit, so that his stopping value must remain positive as well.

Considering the class of games with "experts" and "amateurs" is illustrative, but the result of Theorem 1 is more general. For example, we can apply the same logic to a game in which amateurs receive weak informative signals. The same argument then goes through as long as their stopping value remains positive in the limit for all realizations of their signal.

Let us now contrast the result with the symmetric equilibrium of a game in which agents receive identically distributed signals. In this case, there is no persistent gap in posterior beliefs about the state, and the stopping rate of all agents converges to zero in the limit as $t \to \infty$. Consequently, stopping occurs if and only if one agent's signal exceeds the equilibrium threshold s^* given in Proposition 1.

Theorem 2. *If all agents' signals are identically distributed, then in the unique full-participation equilibirum, there is too little stopping relative to the social optimum.*

What is driving this difference is the information gap between well and poorly informed agents. The stopping time of well informed agents is more strongly correlated with the state of the world, than that of agents whose signal is noisy. This means in particular that when the state is low, then poorly informed agents are less cautious and crowd out participation of the better informed. As a result, the poorly informed learn little from the better informed agents' inaction. In the limit, agents with more accurate information become entirely passive, and the poorly informed effectively play against each other.

6 Conclusion

The objective of this paper was to reveal some of the mechanisms that underly volunteering environments in which information is dispersed and in which learning occurs through different interfering channels. Here, agents learn through private signals and from other's behavior. The basic setup of the model has been kept purposefully simple to retain tractability. It is however natural to consider extensions. For example, first-mover advantage or second-mover advantage appear plausible in many applications, such as R&D competition. Such a change would create a bias among agents for action or inaction, depending on whether we consider first or second-mover advantages, but qualitatively the basic insights in this paper remain the same. Another possibility would be to study how private information affects free-riding in richer model in which experimentation occurs over time contemporaneously with learning from others' action. This alternative was in fact the starting point of a previous version of this paper. However, such a model is significantly more difficult to analyze. Alternatively, 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. Finally, private information arrives here only at the first instant. It would perhaps be feasible to consider a model in which information arrives over time.

A Proofs

Proof of Lemma 1. (i) *In a full-participation equilibrium, strategies are strictly decreasing.* The proof proceeds in two steps. We first show that strategies are non-decreasing as a consequence of supermodularity of the agents' payoff functions. Second, we show that strategies do not have atoms, i.e., for no agent i is there a set of types of positive measure with the property that all types in the set choose the same action t, so that $\Pr(\sigma_i^{-1}(t)) > 0$. These two result together imply that σ_i must be strictly decreasing.

(i.1) *Strategies are non-increasing.* Let $p(s_i) = \Pr(H|s_i)$. The payoff of stopping at time t for agent i with signal s_i is

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

where

$$v_i(t, s_i) = p(s_i)(1 - G_{i,H}(t))H - (1 - p(s_i))(1 - G_{i,L}(t)).$$

Let $\Delta u_i(t, t', s_i) = u_i(t, s_i) - u_i(t', s_i)$ and $\Delta v(t, s_i', s_i) = v_i(t, s_i') - v_i(t, s_i)$. Then, for $t' > t$ and $s_i' > s_i$, we have

$$\begin{aligned}
&\Delta u_i(t, t', s_i') - \Delta u_i(t, t', s_i) \\
&= -(p(s_i') - p(s_i)) \int_t^{t'} e^{-rz} dG_{i,H}(z)\, H + e^{-rt} \Delta v(t, s_i', s_i) - e^{-rt'} \Delta v(t', s_i', s_i) \\
&\geq -(p(s_i') - p(s_i)) e^{-rt'} (G_{i,H}(t') - G_{i,H}(t))H + e^{-rt'} (\Delta v(t, s_i', s_i) - \Delta v(t', s_i', s_i)) \\
&= e^{-rt} \Delta v(t, s_i', s_i) + (p(s_i') - p(s_i))[(G_{i,H}(t') - G_{i,H}(t))H + (G_{i,L}(t') - G_{i,L}(t))] \\
&\geq 0
\end{aligned}$$

Thus u_i is supermodular, so that by Topkis' Monotonicity Theorem we have

$$\sigma_i(s_i) = \arg\max_t u_i(t, s_i)$$

is non-increasing in s_i.

(i.2) *Strategies are atomless.* Suppose to the contrary that there is a set of types for agent i of positive measure with the property that all types in the set choose the same action $t > 0$, so that $\Pr(\sigma_i^{-1}(t)) > 0$. Then for no agent j can it be a best response to stop at time t, ruling out that the equilibrium is a full-participation equilibrium.

(ii) *Strategies are differentiable.* We show that if $(\sigma_i)_{i \in N}$ is a full-participation equilibrium and ϕ_i is the inverse of σ_i, then $F_{i,L}(\phi_i(t))$ is Lipschitz-continuous in t for each i. Because $F_{i,H}$ has full support by hypothesis, it follows then that σ_i is absolutely continuous, and thus differentiable almost everywhere.

Fix agent i with signal $\phi_i(t)$. Because agent i prefers t over $t' > t$, we have $\Delta u_i(t, t', s_i) \geq 0$. On the other hand,

$$\begin{aligned}
\Delta u_j(t, t', s_i) &= p(s_i) \int_0^t e^{-rz} dG_{i,H}(z)\, H + e^{-rt} v_i(t, s_i) \\
&\qquad - p(s_i) \int_0^{t'} e^{-rz} dG_{i,H}(z)\, H - e^{-rt'} v_i(t', s_i) \\
&= e^{-rt} v_i(t, s_i) - e^{-rt'} v_i(t', s_i) \; - p(s_i) \int_t^{t'} e^{-rz} dG_{i,H}(z)\, H
\end{aligned}$$

Thus, it follows

$$e^{-rt} v_i(t, s_i) - e^{-rt'} v_i(t', s_i) \; \geq p(s_i) \int_t^{t'} e^{-rz} dG_{i,H}(z)\, H. \tag{12}$$

We further have

$$\int_t^{t'} e^{-rz} dG_i(z) \; \geq e^{-rt'} \int_t^{t'} dG_i(z) \; = (G_{i,H}(t') - G_{i,H}(t)) e^{-rt'}. \tag{13}$$

Using a zero-addition, we find that

$$
\begin{aligned}
&e^{-rt}v_i(t,s_i) - e^{-rt'}v_i(t',s_i) \\
&= e^{-rt}v_i(t,s_i) - e^{-rt'}v_i(t,s_i) + e^{-rt'}v_i(t,s_i) - e^{-rt'}v_i(t',s_i) \\
&= (e^{-rt} - e^{-rt'})v_i(t,s_i) + e^{-rt'}p(s_i)(G_{i,H}(t') - G_{i,H}(t))H \\
&\qquad - e^{-rt'}(1-p(s_i))(G_{i,L}(t') - G_{i,L}(t)).
\end{aligned} \tag{14}
$$

Now, combining (12) through (14), we obtain

$$
(e^{-rt} - e^{-rt'})v_i(t,s_i) \geq e^{-rt'}(1-p(s_i))(G_{i,L}(t') - G_{i,L}(t))
$$

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

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

where

$$
L(t,t') = \frac{rv_i(t,s_i)}{e^{-rt'}(1-p(s_i))}.
$$

For any $I = [t_0, t_1] \subset \mathbb{R}_+$, $L(t,t')$ is continuous on $I \times I$ and bounded, and therefore $L^* = \max_{(t,t') \in I \times I} L(t,t')$ 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,

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

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

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

The last inequality shows that ϕ_j is locally Lipschitz-continuous for each j and thus by Rademacher's Theorem, it is differentiable almost everywhere on $\mathbb{R}_+$. That ϕ_j is differentiable everywhere follows from the Fundamental Theorem of Calculus, by integrating (7). □

Proof of Lemma 2. Define

$$X_i(s_1, \ldots, s_n) = -\frac{r}{\lambda_{i,\theta}(s_i)}\left(\alpha(s)\left[\frac{1}{n-1}\sum_{j=1}^{n} h_j(s_j) - h_i(s_i)\right]H - 1\right).$$

$F_{i,\theta}$ is differentiable and has full support, and σ_i is differentiable by Lemma 1, and therefore α, $1/\lambda_{i,\theta}$ and each h_i are all differentiable with respect to s_i for $i = 1, \ldots, n$. Thus, by Theorem 7.6.1 in Greenberg (2014), there exists a unique solution to (7). □

Proof of Proposition 1. Let ϕ be the inverse of σ, restricted to the interval $(s^*, \bar{s}]$ and let ϕ_0 denote the inverse of the (implicitly defined) strategy for amateurs. We proof the result in two steps. We show that (i) ϕ and ϕ_0 solve the system of equations (7) and (ii) they are strictly decreasing.

(i) Because of symmetry among experts, we can omit the index i, and use the fact that for every amateur the reverse hazard rate ratio is equal to one. Therefore, the differential equation (7) for each amateur becomes

$$-\phi'(t) = \frac{\alpha_m(\phi(t))\big[n-m-(n-m-1)h\big(\phi(t)\big)\big]H-1}{(n-1)\lambda_L(\phi(t))/r}. \tag{15}$$

By the inverse function theorem, we have $\phi'(t) = 1/\sigma'(\phi(t))$. If we then set $t = \sigma(s)$ and substitute, we obtain a differential equation

$$\sigma'(s) = \frac{(n-1)\lambda_L(s)/r}{\alpha_m(s)\,[n-m-(n-m-1)h(s)]\,H-1},$$

The experts' strategy shown in the theorem is then obtained by integrating with respect to s. The boundary condition $\sigma(\bar{s}) = 0$ implies that the constant of integration is zero. For amateurs, notice that if ϕ_0 is the inverse of their strategy, and $\lambda_{L,0}$ the reverse hazard rate of their signal, then their stopping rate is equal to $-\lambda_{L,0}(\phi(t))\phi_0'(t)$, so that from (10) we obtain

$$\phi_0'(t) = -\frac{1}{\lambda_{i,L}(\phi_0(t))}\frac{r}{n-1}\Big(\alpha_m(\phi(t))\big[(h(\phi(t))-1)m+1\big]H-1\Big)$$

which is identical to the right-hand side of (7). Thus, by construction, the strategies of experts and amateurs solve (7).

(ii) The experts' strategy σ is clearly strictly decreasing and therefore ϕ is as well. For amateurs we show that their stopping rate is non-negative. First, note that since σ is strictly decreasing, $\alpha_m(\phi(t))$ is strictly decreasing in t. Moreover, the reverse hazard rate ratio for the experts' signal distribution is strictly greater than 1, and thus its infimum, denoted by h^*, is greater than 1 as well. Hence, letting $\alpha_m^* = \lim_{t\to\infty}\alpha_m(\phi(t))$, the stopping rate of amateurs at any $t \in [0,\infty)$ is larger than

$$\frac{r}{n-1}\Big(\alpha_m^*\big((h^*-1)m+1\big)H-1\Big). \tag{16}$$

From the equilibrium threshold defined in (9), it follows that

$$1 \leq \alpha_m^*(n - m - (n - m - 1)h^*)H.$$

If we now substitute this inequality into (11), we find that the stopping rate for amateurs has the lower bound

$$\alpha_m^*(h^* - 1)H > 0.$$

□

Proof of Proposition 2. Because of the monotone likelihood ratio property, expected payoffs are non-decreasing in signals. Therefore, if it is optimal to stop for a given signal s_i of some agent i, then it must also be optimal and thus the stopping region is indeed characterized by a profile of thresholds $\hat{s}$. The optimal threshold profile solves

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

The associated Lagrangian is

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

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

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

together with the Kuhn-Tucker conditions $\rho_i(s_i - \underline{s}_i) = 0$ and $\mu_i(\bar{s}_i - s_i) = 0$ and $\rho_i, \mu_i \geq 0$ forall $i \in N$.

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

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

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

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

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

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

for each i.

□

Proof of Theorem 1. Suppose $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a full-participation equilibrium and let ϕ_i be the inverse of σ_i for all $i \in N$. Define $\gamma_i(t) = -\phi'_i(t)/\lambda_{i,L}(\phi_i(t))$ to be the stopping rate of agent i at time t in state L. Let $a_i = \lim_{t\to\infty} \phi_i(t)$ for each $i \in N$ and set $a = (a_1, \ldots, a_n)$. Because of aggregate uncertainty (Assumption 2) there is a subset $M \subseteq N$ so that $a_i > \underline{s}_i$ for all $i \in M$, and $a_i = \underline{s}_i$ for all $i \in N \setminus M$. We show that $M \subset N$.

Since $[0, \infty)$ is the image of σ_i by the assumption of full-participation, and since S_i is bounded for each i, it must be the case that $\lim_{s_i \downarrow a_i} \sigma'_i(s_i) = \infty$ for each $i \in N$. The inverse function theorem then implies $\lim_{t\to\infty} \phi'_i(t) = 0$. For each $i \in M$, we have $\lim_{t\to\infty} \lambda_{i,L}(\phi_i(t)) = \lambda_{i,L}(a_i) > 0$ and therefore $\lim_{t\to\infty} \gamma_i(t) = 0$ for each $i \in M$.

By strong asymmetry in the distribution of information, there exist $i, j \in$

M with $h_i(a_i) > h_j(a_j)$. Then

$$(17)\quad \lim_{t\to\infty} \gamma_i(t) = r\alpha(a)\left[\frac{1}{n-1}\sum_{k=1}^{n} h_k(a_k) - h_i(a_i)\right] H - \frac{r}{n-1}$$
$$< r\alpha(a)\left[\frac{1}{n-1}\sum_{k=1}^{n} h_k(a_k) - h_j(a_j)\right] H - \frac{r}{n-1} = \lim_{t\to\infty} \gamma_j(t).$$

and thus $\lim_{t\to\infty} \gamma_i(t) > 0$ implying $j \notin M$. □

Proof of Theorem 2. By Proposition 1, if all agents' signals are drawn from the same signal distribution, then each agent stops in finite time if and only if his signal is higher than s^*, where s^* solves

$$\alpha(s^*, \ldots, s^*)h(s^*) = 1/H.$$

By Proposition 2, the efficient stopping region is $\{(s_1, \ldots, s_n)|s_i > \hat{s}\}$, where $\hat{s}$ solves

$$\alpha(\hat{s}, \ldots, \hat{s})h(\hat{s}) = 1/nH.$$

Now, $\alpha(s, \ldots, s)h(s)$ is increasing in s by MLRP and hence $\hat{s} < s^*$ implying that there is insufficient stopping in equilibrium. □

References

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

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

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

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

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

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

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

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

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

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

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

Greenberg, M. (2014). *Ordinary Differential Equations.* Wiley.

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

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

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

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

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

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

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

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

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