



Helsinki
Center
of
Economic
Research

Discussion Papers

Delay and Information Aggregation in Stopping Games with Private Information

Pauli Murto

Helsinki School of Economics and HECER

and

Juuso Välimäki

Helsinki School of Economics and HECER

Discussion Paper No. 265

June 2009

ISSN 1795-0562

Delay and Information Aggregation in Stopping Games with Private Information

Abstract

We consider a timing game with private information about a common values payoff parameter. Information is only transmitted through the stopping decisions and therefore the model is one of observational learning. We characterize the symmetric equilibria of the game and we show that even in large games where pooled information is sufficiently accurate for first best decisions, aggregate randomness in outcomes persists. Furthermore, the best symmetric equilibrium induces delay relative to the first best.

JEL Classification: C73, D81, D82, D83

Keywords: optimal stopping, dynamic games, information aggregation.

Pauli Murto

Department of Economics,
Helsinki School of Economics,
P.O.Box 1210, FIN-00101
Finland

e-mail: pauli.murto@hse.fi

Juuso Välimäki

Department of Economics,
Helsinki School of Economics,
P.O.Box 1210, FIN-00101
Finland

e-mail: juuso.valimaki@hse.fi

1 Introduction

We analyze a game of timing where the players are privately informed about the optimal time to stop the game. The stopping decision may, for example, relate to irreversible investment, which is the case analyzed in the real options literature. Our point of departure from that literature is in the nature of uncertainty. Rather than assuming exogenous uncertainty in a publicly observable payoff parameter such as the market price, we consider the case of dispersed private information on the common profitability of the investment. We assume that information is only transmitted through observed actions. In other words, our model is one of observational learning, where communication between players is not allowed.

The key question in our paper is how the individual players balance the benefits from observing other players' actions with the costs of delaying their stopping decision beyond what is optimal based on their own private information. Observational learning is potentially socially valuable, because it allows common values information to spread across players. However, when choosing their optimal timing decisions, the players disregard the informational benefits that their decisions have for the other players. This informational externality leads to too late stopping decisions from the perspective of effective information transmission, and this delay dissipates most of the potential informational benefit from the observational learning. Our main findings are: i) The most informative symmetric equilibrium of the game involves delay, ii) the delay persists even if the number of players is large, iii) information aggregates in random bursts of action, and iv) aggregate uncertainty remains even when aggregate information in the model is sufficiently accurate to determine the optimal investment time.

In our model, the first-best time to invest is common to all players and depends on a single state variable ω . Without loss of generality, we identify ω directly as the first-best optimal time to invest. Since all players have information on ω , their observed actions contain valuable information as long as the actions depend on the players' private information.

The informational setting of the game is otherwise standard for social learning models: The players' private signals are assumed to be conditionally i.i.d. given ω and to satisfy the monotone likelihood property. The payoffs are assumed to be quasi-supermodular in ω and the stopping time t . Given these assumptions, the equilibria in our game are in monotone strategies such that a higher signal implies a later investment decision. Our main characterization result describes a simple way to calculate the optimal stopping moment for each player in the most informative symmetric equilibrium of the game. The

optimal investment time is exactly the optimal moment calculated based on the knowledge that other (active) players are not stopping. The game has also less informative equilibria, where all the players, irrespective of their signals, stop immediately because other players stop as well. These equilibria bear some resemblance to the non-informative equilibria in voting games with common values as in Feddersen & Pesendorfer (1997), and also herding equilibria in the literature on observational learning as in Smith & Sorensen (2000).

In order to avoid complicated limiting procedures, we model the stopping game directly as a continuous-time model with multiple stages. Each stage is to be understood as the time interval between two consecutive stopping actions. At the beginning of each stage, all remaining players choose the time to stop, and the stage ends at the first of these stopping times. The stopping time and the identity of the player(s) to stop are publicly observed, and the remaining players update their beliefs with this new information and start immediately the next stage. This gives us a dynamic recursive game with finitely many stages (since the number of players is finite). Since the stage game strategies are simply functions from the type space to non-negative real numbers, the game and its payoffs are well defined.

The most informative equilibrium path involves two qualitatively different phases. When a stage lasts for a strictly positive amount of time, we say that the game is in the *waiting phase*. Since the equilibrium strategies are monotonic in signals, the fact that no players are currently stopping implies that their signals must be above some cutoff level. This in turn implies that it is more likely that the true state is higher, i.e. the first-best optimal stopping time is later. Thus, during the waiting phase all players update their beliefs gradually upwards. Eventually the waiting phase comes to an end as some player stops. At that point, the remaining players learn that the signal of the stopping player is the lowest possible consistent with equilibrium play, and by monotone likelihood ratio property they update their belief about the state discretely downwards. As a result, a positive measure of player types will find it optimal to stop immediately. If such players exist, the following stage ends at time zero, and the game moves immediately to the next stage, where again a positive measure of types stop at time zero. As long as there are consecutive stages that end at time zero, we say that the game is in the *stopping phase*. This phase ends when the game reaches a stage where no player stops immediately. The game alternates between these two phases until all players have stopped. Notice that information accumulates in an asymmetric manner. Positive information (low signals indicating early optimal action) arrives in quick bursts, while pessimistic information indicating higher signals and the need to delay accumulates gradually.

To understand the source of delay in our model, it is useful to point out an inherent

asymmetry in learning in stopping games. That is, while the players can always revise their stopping decisions forward in time in response to new information, they can not go backward in time if they learn to be too late. In equilibrium, every player stops at the optimal time based on her information at the time of stopping. As a consequence, if at *any* moment during the game the current estimate of the stopping player is too high in comparison to the true state realization, then all the remaining players end up stopping too late. In contrast, errors in the direction of too early stopping times tend to be corrected as new information becomes available.

We obtain the sharpest results for games with a large number of players. First, in the large game limit, almost all the players stop too late relative to the first-best stopping time (except in the case where the state is the highest possible and the first-best stopping time is the last admissible stopping time). The intuition for this result is straight-forward. With a large number of players the pooled information contained in the players' signals is precise. If a non-negligible fraction of players were to stop too early, this would reveal the true state. But then it would be optimal for all players to delay, and this would contradict the presumption of too early stopping. Second, we show that almost all players stop at the same instant of real time (even though they may stop in different stages) where the game also ends. This is because in the informative equilibrium, all observed stopping decisions are informative. With a large number of players, most of the players thus have precise information about state when they stop. But as explained above, information can not aggregate before first-best stopping time, which means that players become aware of the true state too late. This leads to a collapse where all the remaining players stop together fully aware of being too late. Finally, we show that even if we condition on the true state, the time at which the players stop remains stochastic.

Our paper is related to the literature on herding. The paper closest to ours is the model of entry by Chamley & Gale (1994).¹ The main difference to our paper is that in that model it is either optimal to invest immediately or never. We allow a more general payoff structure that allows the state of nature to determine the optimal timing to invest, but which also captures Chamley & Gale (1994) as a special case. This turns out to be important for the model properties. With the payoff structure used in Chamley & Gale (1994), uncertainty is resolved immediately but incompletely at the start of the game. In contrast, our model features gradual information aggregation over time. The information revelation in our model is closely related to our previous paper Murto & Välimäki (2009). In that paper learning over time generates dispersed information about the optimal stop-

¹See also Levin & J.Peck (2008), which extends such a model to allow private information on opportunity costs.

ping point, and information aggregates in sudden bursts of action. Moscarini & Squintani (2008) analyze a R&D race, where the inference of common values information is similar to our model, but as their focus is on the interplay between informational and payoff externalities, they have only two players. Our focus, in contrast, is on the aggregation of information that is dispersed within a potentially large population.

It is also instructive to contrast the information aggregation results in our context with those in the auctions literature. In a k^{th} price auction with common values, Pesendorfer & Swinkels (1997) show that information aggregates efficiently as the number of object grows with the number of bidders. Kremer (2002) further analyzes informational properties of large common values auctions of various forms. In our model, in contrast, the only link between the players is through the informational externality, and that is not enough to eliminate the inefficiencies. The persistent delay in our model indicates failure of information aggregation even for large economies. On the other hand, Bulow & Klemperer (1994) analyzes an auction model that features "frenzies" that resemble our bursts of actions. In Bulow & Klemperer (1994) those are generated by direct payoff externalities arising from scarcity, while in our case they are purely informational.

The paper is structured as follows. Section 2 introduces the basic model. Section 3 establishes the existence of a symmetric equilibrium. Section 4 discusses the properties of the game with a large number of players. In section 5 we illustrate the model by Monte-Carlo simulations. Section 6 concludes with a comparison of our results to the most closely related literature.

2 Model

We consider an N -player game where the players choose optimally when to stop. Denote the set of players by $\mathcal{N} \equiv \{1, \dots, N\}$. The payoff of each player i depends on her stopping time t_i and a random variable ω whose value is initially uncertain to all players and whose prior distribution is $P(\omega)$ on Ω . We take $\Omega = \{\underline{\omega}, \dots, \bar{\omega}\} \subset \mathbb{R} \cup \infty$ to be a finite set. Because we allow for the possibility that $\omega = \infty$, it is natural to allow the actions to be taken in the same set, $T_i = \mathbb{R} \cup \infty$.

The stopping decision is irreversible and yields a payoff

$$v_i(t_i, \omega) = v(t_i, \omega)$$

to player i if she stops at t_i and the state of the world is ω . Notice that we assume symmetric payoffs. Furthermore, we assume that for any fixed ω , $v(t_i, \omega)$ is maximized

at $t_i = \omega$. This allows us to identify ω as the full information optimal stopping time in the game.

We assume also that the payoff function v is quasi-supermodular in t_i, ω :

Assumption 1

$$v_i(t_i, \omega) - v_i(t'_i, \omega)$$

is strictly single crossing in ω and

$$v_i(t_i, \omega') - v_i(t_i, \omega)$$

is strictly single crossing in t_i .

The purpose of this assumption is to guarantee monotonicity of the optimal stopping decisions in additional information. Many examples satisfy this assumption:

- Quadratic loss:

$$v_i(t_i, \omega) = -\beta(t_i - \omega)^2.$$

- Discounted loss from early stopping²:

$$v_i(t_i, \omega) = e^{-r \max\{\omega, t_i\}} \cdot V - e^{-rt_i} \cdot C.$$

Because we allow the players to stop at infinity, we use the topology generated by the one-point compactification of $\mathbb{R} \cup \infty$. We assume throughout that $v(t_i, \omega)$ is continuous in t_i in this topology.³ Under this assumption, individual stopping problems have maximal solutions.

Players are initially privately informed about ω . Player i observes privately a signal θ_i from a joint distribution $G(\theta, \omega)$ on $[\underline{\theta}, \bar{\theta}] \times \Omega$. We assume that the distribution is symmetric across i , and that signals are conditionally i.i.d. Furthermore, we assume that the conditional distributions $G(\theta | \omega)$ and corresponding densities $g(\theta | \omega)$ are well defined and they have full supports $[\underline{\theta}, \bar{\theta}]$ independent of ω .

We assume that the signals satisfy monotone likelihood property (MLRP).

Assumption 2 For all i , $\theta' > \theta$, and $\omega' > \omega$,

$$\frac{g(\theta' | \omega')}{g(\theta | \omega')} > \frac{g(\theta' | \omega)}{g(\theta | \omega)}.$$

²For example the investment model of Chamley & Gale (1994) uses this formulation.

³This assumption holds e.g. under bounded payoff functions and discounting.

This assumption allows us to conclude that optimal individual stopping times for player i , $t_i(\theta)$ is monotonic in other players' types: For all j ,

$$\frac{\partial t_i(\theta)}{\partial \theta_j} \geq 0.$$

Assumption 2 also guarantees that the pooled information in the game becomes arbitrarily revealing of the state as N is increased towards infinity.

Furthermore we make the assumption that the information content in individual signals is bounded.

Assumption 3 *There is a constant $\kappa > 0$ such that*

$$\forall \theta, \omega, \quad \frac{1}{\kappa} > g(\theta, \omega) > \kappa.$$

Finally, we assume that signal densities are continuous in θ :

Assumption 4 *For all ω , $g(\theta | \omega)$ is continuous in θ within $[\underline{\theta}, \bar{\theta}]$.*

2.1 Information and Strategies

We model the stage game $\Gamma(s^k)$, $k = 0, 1, \dots$, as a simultaneous move game where all active players choose a time to stop the game. The (informational) state variable $s^k \in S$ contains all information available to the players at the beginning of stage k , as will be specified shortly. A strategy for player i is a sequence $\tau_i = \{\tau_i^k\}$, where

$$\tau_i^k : [\underline{\theta}, \bar{\theta}] \times S \rightarrow \mathbb{R}_+ \cup \infty.$$

For notational convenience, we suppress the dependence on the public state variable and use notation $\tau_i^k(\theta_i)$ to denote the stopping time of player i in stage k . Players are active if they have not stopped the game in any previous stage. Stage k ends at random time

$$t^k = \min_{i \in \mathcal{N}^k} \tau_i^k,$$

where $\mathcal{N}^k \equiv \mathcal{N} \setminus Q^k$ is the set of active players and Q^k is the set of players that have stopped by stage k :

$$\begin{aligned} Q^0 &= \emptyset, \\ Q^k &= Q^{k-1} \cup \arg \min_{i \in \mathcal{N} \setminus Q^{k-1}} \tau_i^{k-1}(\theta_i). \end{aligned}$$

We denote by N^k the number of active players at the beginning of stage k . The total real time that has elapsed at the beginning of stage k is

$$T^k = \sum_{l=0}^{k-1} t^l.$$

A strategy profile is denoted by $\tau = \{\tau^k\} = \{(\tau_i^k, \tau_{-i}^k)\}$. Within each stage k , the time stopping time t^k and the identities of the players that stop at t^k are public information. Stopping times $\tau_i^k(\theta_i) > t^k$ are not observable to players other than i . This restriction captures our modeling assumption that learning about other players' types is observational. Let

$$\begin{aligned} \mathbf{C}_i^k &\equiv \{ \theta \in [\underline{\theta}, \bar{\theta}] \mid \tau_i^k(\theta) > t^k \}, \\ \mathbf{X}_i^k &\equiv \{ \theta \in [\underline{\theta}, \bar{\theta}] \mid \tau_i^k(\theta) = t^k \}. \end{aligned}$$

Thus, if $i \in \mathcal{N}^k$ stops at stage k , then other players learn that $\theta_i \in \mathbf{X}_i^k$, otherwise they learn that $\theta_i \in \mathbf{C}_i^k$. The state variable includes all information about the players at the beginning of stage k . For each $i \in Q^k$, let $l_i < k$ denote the stage at which i stopped. The state of player i is:

$$s_i^k = \begin{cases} \bigcap_{k' < l_i} \mathbf{C}_i^{k'} \bigcap \mathbf{X}_i^{l_i} & \text{for all } i \in Q^k \\ \bigcap_{k' < k} \mathbf{C}_i^{k'} & \text{for all } i \in \mathcal{N}^k \end{cases}.$$

Hence, s_i^k is a subset of $[\underline{\theta}, \bar{\theta}]$ that contains those signal values that are consistent with observed behavior of i . The state variable $s^k = [s_1^k, \dots, s_N^k]$ contains the available information on all players.

We are interested in symmetric equilibria. We shall show that in such equilibria, players stop in the order of their signals:

$$\tau^k(\theta_i) \leq \tau^k(\theta_j) \text{ if } \theta_i \leq \theta_j \text{ for all } k.$$

We call strategies with this property monotonic. Those strategies have the property that all the inverse images of stopping times are intervals of the type space.

For symmetric monotonic strategies, we can express the state variable more concisely. If we let

$$\begin{aligned} \theta_+^k &= \max_{\theta} \{ \theta \mid \tau^k(\theta) = t^k \}, \\ \theta_-^k &= \min_{\theta} \{ \theta \mid \tau^k(\theta) = t^k \}, \end{aligned}$$

we have:

$$s_i^k = (\theta_+^{k-1}, \bar{\theta}] \text{ for } i \in \mathcal{N}^k.$$

Furthermore,

$$s_i^k = (\theta_-^{l_i}, \theta_+^{l_i}] \text{ for } i \in Q^k.$$

We use notation $\underline{\theta}^k \equiv \theta_+^{k-1}$ to denote the highest type that has stopped before the beginning of stage k . Thus, with symmetric monotonic strategies, it is known at the beginning of stage k that all the remaining players have signals within $(\underline{\theta}^k, \bar{\theta}]$.

2.2 Information during the stage

At the outset of stage k , player i has posterior belief $G(\cdot | s^k, \theta_i)$ on $[\underline{\theta}, \bar{\theta}] \times \Omega$. Since the game remains in stage k only as long as no player stops, the choice of $\tau^k(\theta_i)$ is relevant only as far as

$$\tau^k(\theta_j) \geq \tau^k(\theta_i) \text{ for all } j.$$

With monotonic strategies, conditional on her stopping choice being payoff relevant at instant t in stage k , player i knows that

$$\theta_j \geq \min\{\theta \geq \underline{\theta}^k | \tau^k(\theta) \geq t\} \equiv \underline{\theta}^k(t).$$

As a result, the decision to stop at t must be optimal conditional on this information. To include the information flowing during a stage, we introduce the state variable:

$$s^k(t) = (s_i^k(t)) = \begin{cases} s_i^k & \text{for } i \in Q^k, \\ s_i^k \cap [\underline{\theta}^k(t), \bar{\theta}] & \text{for } i \in \mathcal{N}^k. \end{cases}$$

Notice that a strategy profile $\tau^k = (\tau_1^k, \dots, \tau_N^k)$ induces a distribution of state variables for the next stage. We denote this distribution by $F(s^{k+1} | s^k, \tau^k)$.

2.3 Payoffs

With our recursive definition of the game, the payoffs of each stage game are relatively easy to describe. As long as other players adopt symmetric strategies, player i gets payoff

$$\begin{aligned} V^k(s^k, \mathbf{t}_i, \boldsymbol{\tau}(\theta)) &= \Pr\{t_i^k \leq \min_j \tau^k(\theta_j)\} \mathbb{E}_{G(\cdot | s^k(t_i^k), \theta_i)} v(T^k + t_i^k, \omega) \\ &\quad + \mathbb{E}_{F(s^{k+1} | s^k, \tau^k, t_i^k > \min_j \tau^k(\theta_j))} V^{k+1}(s^{k+1}, \mathbf{t}_i, \boldsymbol{\tau}(\theta)) \end{aligned}$$

from strategy $\mathbf{t}_i = \{t_i^k\}$ when other players play according to $\boldsymbol{\tau}(\theta) = \{\tau^k(\theta)\}$. The first expectation on the right hand side is taken with respect to the posterior on the state ω conditional on own stopping. The second expectation on the continuation payoff from stage $k+1$ onwards given the information that i was not the first to stop.

3 Symmetric equilibria

We shall show in this section that the game has always a symmetric equilibrium that we call the *informative equilibrium*. We also note that in some stages the game may have another equilibrium, where all players stop at the beginning of the stage irrespective of their signals. We call such equilibria *uninformative*.

All symmetric equilibria of the game have the property that each player stops at the first moment that is the optimal stopping time conditional on the information received so far, under the extra assumption that *no further information will ever be obtained from other players*. This myopia property makes the computation of the equilibrium straightforward.

3.1 Informative equilibrium

Consider the beginning of an arbitrary stage k , where set of players that have not yet stopped is given by \mathcal{N}^k and it is common knowledge that all of them have signals within $(\underline{\theta}^k, \bar{\theta}]$, where the lowest possible type is given by $\underline{\theta}^k \equiv \theta_+^{k-1}$. To define the informative equilibrium, it is useful to introduce an auxiliary state variable $s^k(\theta) = [s_1^k(\theta), \dots, s_N^k(\theta)]$, where:

$$s_i^k(\theta) \equiv \begin{cases} s_i^k & \text{for } i \in Q^k \\ [\underline{\theta}, \bar{\theta}] & \text{for } i \in N \setminus Q^k \end{cases} .$$

This state variable has the following meaning: A player of type θ that conditions on state $s^k(\theta)$ assumes that all the remaining players have a signal higher than θ . Let us define $\tau_*^k(\theta)$ as the optimal stopping time for such a player:

$$\tau_*^k(\theta) \equiv \inf \{ t \geq 0 \mid \mathbb{E}[v(t, \omega) \mid s^k(\theta)] \geq \mathbb{E}[v(t', \omega) \mid s^k(\theta)] \text{ for all } t' \geq t \} . \quad (1)$$

Note that (1) allows $\tau_*^k(\theta) = \infty$. The following Lemma states that $\tau_*^k(\theta)$ is increasing in θ (strictly so when $0 < \tau_*^k(\theta) < \infty$), and therefore defines a symmetric monotonic strategy profile:

Lemma 1 (Monotonicity of $\tau_*^k(\theta)$) *Let $\tau_*^k(\theta)$ denote the stopping strategy defined in (1).*

- *If $0 < \tau_*^k(\theta) < \infty$ for some $\theta \in (\underline{\theta}, \bar{\theta})$, then for all $\theta' \in [\underline{\theta}, \theta)$ and $\theta'' \in (\theta, \bar{\theta}]$, we have*

$$\tau_*^k(\theta') < \tau_*^k(\theta) < \tau_*^k(\theta'').$$

- *If $\tau_*^k(\theta) = 0$ for some $\theta \in (\underline{\theta}, \bar{\theta})$, then for all $\theta' \in [\underline{\theta}, \theta)$ we have $\tau_*^k(\theta') = 0$.*

- If $\tau_*^k(\theta) = \infty$ or some $\theta \in (\underline{\theta}, \bar{\theta})$, then for all $\theta'' \in (\theta, \bar{\theta}]$ we have $\tau_*^k(\theta'') = \infty$.

Proof. These follow directly from the Assumptions 1 and 2. ■

The next Theorem states that this profile is an equilibrium. The proof makes use of the one-step deviation principle and the assumption of MLRP. We call this profile the informative equilibrium of the game.

Theorem 1 (Informative equilibrium) *The game has a symmetric equilibrium, where every player adopts at stage k the strategy $\tau_*^k(\theta)$ defined in (1).*

Proof. Assume that all players i use strategies given by (1) in each stage k . It is clear that no player can benefit by deviating to $\tau_i < \tau_*^k(\theta_i)$. Let $\hat{\tau}_i(\theta_i) > \tau_*^k(\theta_i)$ be the best deviation for player i of type θ_i in stage k . Let $\hat{\theta}_i$ be the type of player i that solves

$$\tau_*^k(\hat{\theta}_i) = \hat{\tau}_i(\theta_i).$$

By Assumptions 1 and 2, we know that $\hat{\theta}_i > \theta_i$, and also that

$$\mathbb{E} \left[v(t, \omega) \mid s^k(\hat{\theta}_i), \theta_i \right]$$

is decreasing in t at $t = \tau_*^k(\hat{\theta}_i)$. But this contradicts the optimality of the deviation to $\hat{\tau}_i(\theta_i)$.

Since there are no profitable deviations in a single stage for any type of player i , the claim is proved by the one-shot deviation principle. ■

Let us next turn to the properties of the informative equilibrium. The equilibrium stopping strategy $\tau_*^k(\theta)$ defines a time dependent cutoff level $\theta^k(t)$ for all $t \geq 0$ as follows:

$$\theta^k(t) \equiv \begin{cases} \underline{\theta}^k & \text{if } 0 \leq t < \tau_*^k(\underline{\theta}^k) \\ \bar{\theta} & \text{if } t > \tau_*^k(\bar{\theta}) \\ \max \{ \theta \mid \tau_*^k(\theta) \leq t \} & \text{if } \tau_*^k(\underline{\theta}^k) \leq t \leq \tau_*^k(\bar{\theta}) \end{cases}. \quad (2)$$

In words, $\theta^k(t)$ is the highest type that stops at time t in equilibrium. The key properties of $\theta^k(t)$ for the characterization of equilibrium are given in Proposition 1 below. Before that, we note that the equilibrium stopping strategy is left-continuous in θ :

Lemma 2 (Left-continuity of $\tau_*^k(\theta)$) *Let $\tau_*^k(\theta)$ denote the informative equilibrium stopping strategy defined in (1). For all $\theta \in (\underline{\theta}, \bar{\theta})$,*

$$\lim_{\theta' \uparrow \theta} \tau_*^k(\theta') = \tau_*^k(\theta).$$

Proof. Assume on the contrary that for some θ , we have $\tau_*^k(\theta) - \lim_{\theta' \uparrow \theta} \tau_*^k(\theta') > 0$ (Lemma 1 guarantees that we can not have $\tau_*^k(\theta) - \lim_{\theta' \uparrow \theta} \tau_*^k(\theta') < 0$). Denote $\Delta t = t'' - t'$, where $t'' = \tau_*^k(\theta)$ and $t' = \lim_{\theta' \uparrow \theta} \tau_*^k(\theta')$. Denote $u(t, \theta) = \mathbb{E}[v(t, \omega) | s^k(\theta)]$. By definition of $\tau_*^k(\theta)$, we have then $u(t'', \theta) > u(t, \theta)$ for all $t \in [t' - \delta, t' + \delta]$ for any $0 < \delta < \Delta t$.

Because signal densities are continuous in θ , $u(t, \theta)$ must be continuous in θ . This means that there must be some $\varepsilon > 0$ such that $u(t'', \theta') > u(t, \theta')$ for all $t \in [t' - \delta, t' + \delta]$ and for all $\theta' \in [\theta - \varepsilon, \theta]$. But on the other hand $\lim_{\theta' \uparrow \theta} \tau_*^k(\theta') = t'$ implies that $\tau_*^k(\theta') \in [t' - \delta, t' + \delta]$ if θ' is chosen sufficiently close to θ . By definition of $\tau_*^k(\theta')$ this means that $u(\tau_*^k(\theta')) \geq u(t'', \theta')$, and we have a contradiction. We can conclude that for all θ , $\lim_{\theta' \uparrow \theta} \tau_*^k(\theta') = \tau_*^k(\theta)$. ■

The next proposition allows us to characterize the key properties of the informative equilibrium. It says that $\theta^k(t)$ is continuous, which means that at each $t > 0$, only a single type exists, and hence the probability of more than one player stopping simultaneously is zero for $t > 0$. In addition, the Proposition says that along equilibrium path, $\theta^k(0) > \underline{\theta}^k$ for all stages except possibly the first one. This means that at the beginning of each stage there is a strictly positive probability that many players stop simultaneously.

Proposition 1 $\theta^k(t) : [0, \infty) \rightarrow [\underline{\theta}^k, \bar{\theta}]$ defined in (2) is continuous, (weakly) increasing, and along the path of the informative equilibrium $\theta^k(0) > \underline{\theta}^k$ for $k \geq 1$.

Proof. Continuity and monotonicity of $\theta^k(t)$ follow from definition (2) and the properties of $\tau_*^k(\theta)$ given in Lemmas 1 and 2.

Take any stage $k \geq 1$ along the informative equilibrium path. To see that we must have $\theta^k(0) > \underline{\theta}^k$, consider how information of the marginal player changes at time t^{k-1} . If $t^{k-1} = 0$, the player with signal $\theta_+^{k-1} = \underline{\theta}^k$ was willing to stop at $t^{k-1} = 0$ conditional on being the lowest type within the remaining players. However, since the stage ended at $t^{k-1} = 0$, at least one player had a signal within $(\theta_i^{k-1}, \theta_+^{k-1}]$. By MLRP and quasi-supermodularity, this additional information updates the beliefs of the remaining players discretely downwards. Therefore, $\tau_*^k(\theta) = 0$ for all $\theta \in [\underline{\theta}^k, \theta^k + \varepsilon]$ for some $\varepsilon > 0$, which by (2) means that $\theta^k(0) > \underline{\theta}^k$.

On the other hand, if $t^{k-1} > 0$, the lowest signal within the remaining players in stage $k - 1$ was $\theta_+^{k-1} = \underline{\theta}^k$. The player with this signal stopped optimally under the information that all the remaining players have signals within $[\underline{\theta}^k, \bar{\theta}]$. But as soon this player stops and the game moves to stage k , the other players update on the information that one of the players remaining in the game in stage $k - 1$ had the lowest possible signal

value amongst the remaining players. Again, by MLRP and quasi-supermodularity, the marginal cutoff moves discretely upwards, and we have $\theta^k(0) > \underline{\theta}^k$. ■

To understand the equilibrium dynamics, note that as real time moves forward, the cutoff $\theta^k(t)$ moves upward, thus shrinking from left the interval within which the signals of the remaining players lie. By MLRP and quasi-modularity this new information works towards delaying optimal stopping time for all the remaining players. At the same time, keeping information fixed, the passage of time brings forth the optimal stopping time for additional types. In equilibrium, $\theta^k(t)$ moves at a rate that exactly balances these two effects keeping the marginal type indifferent. As soon as the stage ends at $t^k > 0$, the expected value from staying in the game drops by a discrete amount for the remaining players (again by MLRP and quasi-supermodularity). This means that the marginal cutoff moves discretely upwards and thus $\theta^{k+1}(0) > \theta^k(t^k) = \underline{\theta}^{k+1}$, and at the beginning of the new stage there is thus a mass point of immediate exits. If at least one player stops, the game moves immediately to stage $k + 2$ with another mass point of exits, and this continues as long as there are consecutive stages in which at least one player stops at $t = 0$. Thus, the equilibrium path alternates between "stopping phases", i.e. consecutive stages that end at $t = 0$ and result with multiple simultaneous exits, and "waiting phases", i.e. stages that continue for a strictly positive time.

Note that the random time at which stage k ends,

$$t^k = \tau_*^k \left(\min_{i \in \mathcal{N}^k} \theta_i \right),$$

is directly linked to the first order statistic of the player types remaining in the game at the beginning of stage k . If we had a result stating that for all k , $\tau^k(\theta_i)$ is strictly increasing in θ_i , then the description of the equilibrium path would be equivalent to characterizing the sequence of lowest order statistics where the realization of all previous statistics is known. Unfortunately this is not the case, since for all stages except the very first one there is a strictly positive mass of types that stop immediately at $t = 0$, which means that the signals of those players will be revealed only to the extent that they lie within a given interval. However, in Section 4.3 we will show that in the limit where the number of players is increased towards infinity, learning in equilibrium is equivalent to learning sequentially the exact order statistics of the signals.

3.2 Uninformative equilibria

While the model always admits the existence of the informative symmetric equilibrium defined above, some stage games also allow the possibility of an additional symmetric

equilibrium, where all players stop at the beginning of the stage irrespective of their signals. We call these uninformative equilibria.

To understand when such uninformative equilibria exist, consider the optimal stopping time of a player who has private signal θ , conditions on all information s^k obtained in all stages $k' < k$, but who does not obtain any new information in stage k . Denote the optimal stopping time of such a player by $\bar{\tau}^k(\theta)$:

$$\bar{\tau}^k(\theta) \equiv \min \{t \geq 0 \mid \mathbb{E}[v(t, \omega) \mid s^k, \theta] \geq \mathbb{E}[v(t', \omega) \mid s^k, \theta] \text{ for all } t' \geq t\}.$$

If $\bar{\tau}^k(\theta) > 0$ for some $\theta \in [\underline{\theta}, \bar{\theta}]$, then an uninformative equilibrium can not exist: it is a strictly dominant action for that player to continue beyond $t = 0$. But if $\bar{\tau}^k(\theta) = 0$ for all players, then an uninformative equilibrium indeed exists: If all players stop at $t = 0$ then they learn nothing from each other. And if they learn nothing from each other, then $t = 0$ is their optimal action.

Since $\bar{\tau}^k(\theta)$ is clearly increasing in θ , the existence of uninformative equilibria depends simply on whether $\bar{\tau}^k(\bar{\theta})$ is zero:

Proposition 2 *If at stage k we have $\bar{\tau}^k(\bar{\theta}) = 0$, then the game has a symmetric equilibrium, where at stage k all active players stop at time $\tau^k = 0$ irrespective of their signals.*

The equilibrium, where all the active players choose $\tau^k = 0$ in *all* stages with $\bar{\tau}^k(\bar{\theta}) = 0$, is the *least informative equilibrium* of the game. There are of course also intermediate equilibria between the informative and least informative equilibria, where at some stages with $\bar{\tau}^k(\bar{\theta}) = 0$ players choose $\tau^k(\theta)$ defined in (1), and in others they choose $\tau = 0$.

Note that there are also stages where the informative equilibrium commands all players to stop at $t = 0$. This happens if the remaining players are so much convinced that they have already passed the optimal stopping time that even finding out that all of them have signals $\theta = \bar{\theta}$ would not make them think otherwise. In that case $\tau^k(\theta) = 0$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$, where $\tau^k(\theta)$ is defined in (1).

It is easy to rank the symmetric equilibria of the game. The informative equilibrium is payoff dominant in the class of all symmetric equilibria of the game. The option of stopping the game is always present for all players in the game, and as a result, not stopping must give at least the same payoff.

4 Informative Equilibrium in Large Games

In this section we study the limiting properties of the model, when we increase the number of players towards infinity. In subsection 4.1 we show that the informative equilibrium

exhibits delay and randomness. In subsection 4.2 we discuss the effect on the players' payoffs of the observational learning. In subsection 4.3 we analyze the information of the players in equilibrium, and derive a simple algorithm for simulating the equilibrium path directly in the large game limit.

4.1 Delay in Equilibrium

We state here a theorem that characterizes the equilibrium behavior in the informative equilibrium for the model with a general state space Ω in the limit $N \rightarrow \infty$. Let $T_N(\theta, \omega)$ denote the random exit time (in real time) in the informative equilibrium of a player with signal θ when the state is ω and the number of players at the start of the game is N . We will be particularly interested in the behavior of $T_N(\theta, \omega)$ as N grows and we define

$$T(\omega, \theta) \equiv \lim_{N \rightarrow \infty} T_N(\omega, \theta),$$

where the convergence is to be understood in the sense of weak convergence.⁴ Since we have assumed Ω to be compact, we know that the sequence $T_N(\theta, \omega)$ has a convergent subsequence. For now, we take $T(\omega, \theta)$ to be the limit of any such subsequence. Along the way, we shall prove that this is also the limit of the original sequence.

The real time instant when the last player with signal $\bar{\theta}$ stops is given by $T_N(\omega, \bar{\theta})$ and we let

$$T_N(\omega) \equiv T_N(\omega, \bar{\theta}) \text{ and } T(\omega) \equiv \lim_{N \rightarrow \infty} T_N(\omega).$$

We let $F(t | \omega)$ denote the distribution of $T(\omega)$, or in other words,

$$F(t | \omega) = \Pr\{T(\omega) \leq t\},$$

and use $f(t | \omega)$ to refer to the corresponding probability density function. The following Theorem characterizes the asymptotic behavior of the informative equilibrium as the number of players becomes large.

Theorem 2 *In the informative equilibrium of the game, we have for all $\omega < \bar{\omega}$,*

1. $\text{supp}f(t | \omega) = [\max\{t(\underline{\theta}), \omega\}, \bar{\omega}]$.
2. For all $\theta, \theta' \in (\underline{\theta}, \bar{\theta}]$,

$$\lim_{N \rightarrow \infty} \Pr\{T_N(\omega, \theta) = T_N(\omega, \theta')\} = 1.$$

⁴In our setting, this is also equivalent to convergence in distribution.

Proof. In a symmetric equilibrium, no information is transmitted before the first exit. By monotonicity of the equilibrium strategies, a lower bound for all exit times and hence also for $T_N(\omega)$ for all N is $t(\underline{\theta})$.

Consider next an arbitrary $\theta' > \underline{\theta}$. By the law of large numbers, we have for all ω :

$$\frac{\#\{i \in \{1, \dots, N\} \mid \theta_i < \theta'\}}{N} \rightarrow G(\theta' \mid \omega).$$

By Assumption 3, and the law of large numbers, for each θ' there is a $\theta'' < \theta'$ such that for all $\omega < \bar{\omega}$ and all $t < \omega$

$$\lim_{N \rightarrow \infty} \Pr\{\exists k \text{ such that } \theta_+^k < \theta'' < \theta' < \theta_+^{k+1}\} = 0.$$

This follows from the fact that

$$\lim_{\theta'' \rightarrow 0} \frac{G(\theta'' \mid \omega)}{G(\theta' \mid \omega)} = 0,$$

and the fact that by Assumption 2, for all $\omega' \neq \bar{\omega}$,

$$\lim_{N \rightarrow \infty} \left(\frac{G(\theta' \mid \omega')}{G(\theta' \mid \bar{\omega})} \right)^N = 0.$$

Consider therefore the stage k' where a player with the signal θ' stops. Then $\theta'' < \theta_+^{k'-1} < \theta'$, and the player with signal θ'' knows

$$\frac{\#\{i \in \{1, \dots, N\} \mid \theta_i < \theta_+^{k'-1}\}}{N}.$$

By the law of large numbers, this is sufficient to identify ω . This implies part 2 of the Theorem and also that $\text{supp}f(t \mid \omega) \subset [\max\{t(\underline{\theta}), \omega\}, \bar{\omega}]$.

The lower bound of the support is by the argument above $\max\{t(\underline{\theta}), \omega\}$, and the remaining task is to argue that the upper bound of the support is $\bar{\omega}$. This follows easily from the fact that if $\Pr\{T_N(\omega) < t\} \rightarrow 1$ for some $t < \bar{\omega}$, then the first exit must take place before t with probability 1 but this is inconsistent with symmetric informative equilibrium in monotonic strategies. To see this, let $t' \leq t$ be the smallest instant such that

$$\lim_{N \rightarrow \infty} \Pr\{\exists i \in \{1, \dots, N\} : \tau_*^1(\theta_i) \leq t'\} = 1.$$

By Assumption 2, conditional on no exit by t' , the posterior probability on Ω converges to a point mass on $\bar{\omega}$. ■

4.2 Payoffs in equilibrium

We turn next to the effect of observational learning on the players' payoffs. To be precise about this, we define three ex-ante payoff functions. First, we denote by V^0 the ex-ante value of a player whose belief on the state is given by the prior:

$$V^0 = \sum_{\omega \in \Omega} \pi^0(\omega) v(T^0, \omega),$$

where T^0 is the optimal timing based on the prior only:

$$T^0 = \arg \max_t \sum_{\omega \in \Omega} \pi^0(\omega) v(t, \omega).$$

Second, consider a player who has a private signal but does not observe other players. The ex-ante value of such an "isolated" player is:

$$V^1 = \sum_{\omega \in \Omega} \left[\pi^0(\omega) \int_{\underline{\theta}}^{\bar{\theta}} g(\theta | \omega) v(T^\theta, \omega) d\theta \right],$$

where T^θ is the optimal stopping time with signal θ and $\pi^\theta(\omega)$ is the corresponding posterior:

$$\begin{aligned} T^\theta &\equiv \arg \max_t \sum_{\omega \in \Omega} \pi^\theta(\omega) v(t, \omega), \\ \pi^\theta(\omega) &\equiv \frac{\pi^0(\omega) g(\theta | \omega)}{\sum_{\omega \in \Omega} \pi^0(\omega) g(\theta | \omega)}. \end{aligned}$$

Third, consider a player in the informative equilibrium of the game. We assume that N is very large, which by Theorem 2 means that almost all players stop at the same random time $T(\omega)$ (the moment of collapse). From an ex-ante point of view, the equilibrium payoff is determined by its probability distribution $f(t | \omega)$. The ex-ante equilibrium payoff is thus:

$$V^* = \sum_{\omega \in \Omega} \left[\pi^0(\omega) \int_0^\infty f(t | \omega) v(t, \omega) dt \right]. \quad (3)$$

It is clear that additional learning can never reduce the ex-ante value, and therefore we must have:

$$V^* \geq V^1 \geq V^0.$$

We call $V^P \equiv V^1 - V^0$ the value of private learning, and $V^S \equiv V^1 - V^*$ the value of social learning. In Section 5 we demonstrate numerically that V^S and V^P are closely related to

each other. In particular, the value of social information increases as the value of private information is increased. We can also derive analytically an upper bound for V^S , which shows that whenever the individual private signals are non-informative in the sense that V^P is very small, then also V^S must be small (this holds even if the pooled information is still arbitrarily informative).

An important effect of observational learning is that it increases the sensitivity of players' payoffs to the realized state of nature. We will demonstrate this effect numerically in Section 5. We can also define value functions conditional on realized signal:

$$\begin{aligned} V^1(\theta) &= \sum_{\omega \in \Omega} \pi^\theta(\omega) v(T^\theta, \omega), \\ V^*(\theta) &= \sum_{\omega \in \Omega} [\pi^\theta(\omega) V^*(\omega)]. \end{aligned}$$

We conjecture that $V^S(\theta) \equiv V^*(\theta) - V^1(\theta)$ is increasing in θ , that is, the additional value of observational learning is more valuable to players who have obtained a high signal. The intuition runs as follows. If the true state is low, a player with a high signal benefits a lot from the information released by the other players who have low signals (since they will act before her). But if the true state is high, a player with a low signal will learn nothing from the other players that have higher signals (because those players will act after her). The computations in Section 5 support this conjecture.

It is clear that the player with the lowest possible signal cannot benefit from observational learning at all (she must be indifferent between following her own signal and following an equilibrium strategy), and we must therefore have

$$V^1(\underline{\theta}) = V^*(\underline{\theta}).$$

4.3 Information in equilibrium

The properties of the informative equilibrium rely on the statistical properties of the order statistics of the players' signals. In this subsection we analyze the information content in those order statistics in the limit $N \rightarrow \infty$.

Denote the n :th order statistic in the game with N players by

$$\tilde{\theta}_n^N \equiv \min \{ \theta \in [\underline{\theta}, \bar{\theta}] \mid \# \{ i \in \mathcal{N} \mid \theta_i \leq \theta \} = n \}. \quad (4)$$

It is clear that if we now increase N towards infinity while keeping n fixed, $\tilde{\theta}_n^N$ converges to $\underline{\theta}$ in probability. Therefore, it is more convenient to work with random variable

$$Y_n^N \equiv \left(\tilde{\theta}_n^N - \underline{\theta} \right) \cdot N. \quad (5)$$

Note that Y_n^N has the same information content as $\tilde{\theta}_n^N$, but as we will show below, it will converge in distribution to a non-degenerate random variable. This limit distribution, therefore, captures the information content of $\tilde{\theta}_n^N$ in the limit. Let us also define

$$\Delta Y_n^N \equiv Y_n^N - Y_{n-1}^N = \left(\tilde{\theta}_n^N - \tilde{\theta}_{n-1}^N \right) \cdot N, \quad (6)$$

where by convention we let $\theta_0^N \equiv \underline{\theta}$ and $Y_0^N \equiv 0$. The following proposition shows that ΔY_n^N converge to independent exponentially distributed random variables as $N \rightarrow \infty$:

Proposition 3 Fix $n \in \mathbb{N}_+$ and denote by $[\Delta Y_1^\infty, \dots, \Delta Y_n^\infty]$ a vector of n independent exponentially distributed random variables with parameter $g(\underline{\theta} \mid \omega)$:

$$\Pr(\Delta Y_1^\infty \leq x_1, \dots, \Delta Y_n^\infty \leq x_n) = e^{-g(\underline{\theta} \mid \omega) \cdot x_1} \cdot \dots \cdot e^{-g(\underline{\theta} \mid \omega) \cdot x_n}.$$

Consider the sequence of random variables $\{[\Delta Y_1^N, \dots, \Delta Y_n^N]\}_{N=n}^\infty$, where for each N the random variables ΔY_i^N are defined by (4) - (6). As $N \rightarrow \infty$, we have:

$$[\Delta Y_1^N, \dots, \Delta Y_n^N] \xrightarrow{\mathcal{D}} [\Delta Y_1^\infty, \dots, \Delta Y_n^\infty],$$

where $\xrightarrow{\mathcal{D}}$ denotes convergence in distribution.

Proof. The probability distribution of ΔY_n^N , conditional on Y_{n-1}^N is given by:

$$\begin{aligned} \Pr(\Delta Y_n^N \leq x \mid Y_{n-1}^N) &= \Pr\left(\left(\tilde{\theta}_n^N - \tilde{\theta}_{n-1}^N\right) \cdot N \leq x \mid \tilde{\theta}_{n-1}^N\right) \\ &= \Pr\left(\tilde{\theta}_n^N \leq \tilde{\theta}_{n-1}^N + \frac{x}{N} \mid \tilde{\theta}_{n-1}^N\right) \\ &= 1 - \left(1 - \frac{\left(G\left(\tilde{\theta}_{n-1}^N + \frac{x}{N} \mid \omega\right) - G\left(\tilde{\theta}_{n-1}^N \mid \omega\right)\right)}{1 - G\left(\tilde{\theta}_{n-1}^N \mid \omega\right)}\right)^{N-n}. \end{aligned}$$

Noting that as $N \rightarrow \infty$, we have $\tilde{\theta}_{n-1}^N \xrightarrow{\mathcal{P}} \underline{\theta}$ and $\frac{x}{N} \rightarrow 0$, and therefore we have:

$$N \cdot \frac{\left(G\left(\tilde{\theta}_{n-1}^N + \frac{x}{N} \mid \omega\right) - G\left(\tilde{\theta}_{n-1}^N \mid \omega\right)\right)}{1 - G\left(\tilde{\theta}_{n-1}^N \mid \omega\right)} \xrightarrow{\mathcal{P}} g(\underline{\theta} \mid \omega) \cdot x.$$

Noting also that

$$\lim_{N \rightarrow \infty} \left(1 - \frac{g(\underline{\theta} \mid \omega) \cdot x}{N}\right)^{N-n} = e^{-g(\underline{\theta} \mid \omega) \cdot x},$$

we have:

$$\lim_{N \rightarrow \infty} \Pr(\Delta Y_n^N \leq x \mid Y_{n-1}^N) = 1 - e^{-g(\underline{\theta} \mid \omega) \cdot x}.$$

This means that ΔY_n^N converges in distribution to an exponentially distributed random variable with parameter $g(\underline{\theta} \mid \omega)$ that is independent of all lower order statistics. ■

Note that the limit distribution of ΔY_n^N does not depend on n . Therefore, $Y_n^N = \sum_{i=1}^n \Delta Y_i^N$ converges to a sum of independent exponentially distributed random variables, which means that the limiting distribution of Y_n^N is Gamma distribution:

Corollary 3 Y_n^N converges to a Gamma distributed random variable:

$$Y_n^N = \sum_{i=1}^n \Delta Y_i^N \xrightarrow{\mathcal{D}} \sum_{i=1}^n \Delta Y_i^\infty \equiv Y_n^\infty,$$

where $Y_n^\infty \sim \text{Gamma}(n, g(\underline{\theta} | \omega))$.

Proposition 3 means that when N is large, observing the n lowest order statistics is observationally equivalent to observing n independent exponentially distributed random variables. This has an important implication for the Bayesian updating based on order statistics: observing only the n :th order statistic $\tilde{\theta}_n^N$ is informationally equivalent to observing $\{\tilde{\theta}_i^N\}_{i=1}^n$ that contains all order statistics up to n . This is due to the "memoryless" nature of exponential random variables. To see this formally, write the posterior belief of an observer who updates her belief on the state of the world based on the realization $\{\tilde{\theta}_i^N\}_{i=1}^n$ (approximating the joint distribution of $(\tilde{\theta}_i^N - \tilde{\theta}_{i-1}^N) \cdot N$ by exponential distribution based on Proposition 3). As can be seen, this posterior depends only on the realization of $\tilde{\theta}_n^N$:

$$\begin{aligned} \pi\left(\omega \mid \{\tilde{\theta}_i^N\}_{i=1}^n\right) &\approx \frac{\pi^0(\omega) \cdot \prod_{i=1}^n g(\underline{\theta} | \omega) e^{-g(\underline{\theta} | \omega)(\tilde{\theta}_i^N - \tilde{\theta}_{i-1}^N) \cdot N}}{\sum_{\omega \in \Omega} \pi^0(\omega) \cdot \prod_{i=1}^n g(\underline{\theta} | \omega) e^{-g(\underline{\theta} | \omega)(\tilde{\theta}_i^N - \tilde{\theta}_{i-1}^N) \cdot N}} \\ &= \frac{\pi^0(\omega) \cdot (g(\underline{\theta} | \omega))^n e^{-g(\underline{\theta} | \omega)(\tilde{\theta}_n^N - \underline{\theta}) \cdot N}}{\sum_{\omega \in \Omega} \pi^0(\omega) \cdot (g(\underline{\theta} | \omega))^n e^{-g(\underline{\theta} | \omega)(\tilde{\theta}_n^N - \underline{\theta}) \cdot N}}. \end{aligned}$$

So far, we have discussed the properties of the order statistics of the signals without linking them to the equilibrium behavior. Now we turn to the properties of the informative equilibrium, and show that in the large game limit the equilibrium path can be approximated by a simple algorithm that samples sequentially the order statistics. To make this statement precise, we now fix N and define two different sequences of random variables, both obtained as mappings from realized signal values to real numbers.

First, for each N , denote by $T_n^*(N)$ the real time at which the number of players that stop exceeds n in the unique informative equilibrium:

$$T_n^*(N) \equiv \min \{T^k \mid Q^k \geq n\}.$$

The increasing sequence $\{T_n^*(N)\}_{n=1}^N$ contains the real stopping moments of all N players in the game.

Second, we define a sequence of stopping times $\widehat{T}_n(N)$ calculated directly on the basis of the order statistics. As an intermediate step, denote by $\widetilde{T}_n(N)$ the optimal stopping moment given the information contained in the n lowest order statistics:

$$\widetilde{T}_n(N) \equiv \inf \left\{ t \geq 0 \mid \mathbb{E} \left[v(t, \omega) \mid \left\{ \widetilde{\theta}_i^N \right\}_{i=1}^n \right] \geq \mathbb{E} \left[v(t', \omega) \mid \left\{ \widetilde{\theta}_i^N \right\}_{i=1}^n \right] \text{ for all } t' \geq t \right\}.$$

Next, define random variable $\widehat{T}_n(N)$ as:

$$\widehat{T}_n(N) \equiv \max_{i=1, \dots, n} \widetilde{T}_n(N). \quad (7)$$

Hence, $\{\widehat{T}_n(N)\}_{n=1}^N$ is the sequence of optimal stopping times based on sequential sampling of order statistics under an additional constraint that one is never allowed to "go back in time", i.e. choose a stopping time lower than some previously chosen stopping time. Note that both $\{T_n^*(N)\}_{n=1}^N$ and $\{\widehat{T}_n(N)\}_{n=1}^N$ are weakly increasing sequences of random variables.

The next proposition says that for any fixed n , the difference between $\widehat{T}_n(N)$ and $T_n^*(N)$ vanishes as N goes to infinity (in the sense of convergence in probability). The key for this result is the finding that inference on order statistics becomes informationally equivalent to inference based on independent exponentially distributed random variables. This means that a player that conditions on having the lowest signal among the remaining players does not learn anything more by conditioning on exact realizations of the signals lower than hers. Thus, inference based on the exact realizations of lowest order statistics becomes the same as the inference of the marginal player in equilibrium, who knows the lowest signal realizations only to the extent that they lie within some fixed intervals.

Proposition 4 Fix $n \in \mathbb{N}_+$ and consider random variables $T_n^*(N)$ and $\widehat{T}_n(N)$. As $N \rightarrow \infty$, we have:

$$\widehat{T}_n(N) - T_n^*(N) \xrightarrow{\mathcal{P}} 0.$$

Proof. Fix n . As $N \rightarrow \infty$, the updating based on the realizations of the n lowest signals is informationally equivalent to observing n exponentially distributed random variables with parameter $g(\underline{\theta} \mid \omega)$. Consider the player that has the n :th lowest signal $\widetilde{\theta}_n^N$. As N is increased, this signal is of course arbitrarily close to $\underline{\theta}$ at a probability arbitrarily close to one. In equilibrium, this player is the n :th to stop (possibly together with some other players). By (1), her real stopping time $T_n^*(N)$ is optimal conditional on information that some $n' < n$ players have signals within $[\underline{\theta}, \theta']$ for some $\theta' \leq \widetilde{\theta}_n^N$, no player as signals

within $(\theta', \tilde{\theta}_n^N)$, and she herself has signal $\tilde{\theta}_n^N$. In contrast, $\tilde{T}_n(N)$ is optimal conditional on n players having signals within $[\underline{\theta}, \tilde{\theta}_n^N]$, which by MLRP and super-modularity means that for any $\varepsilon > 0$,

$$\lim_{N \rightarrow \infty} \Pr \left(\tilde{T}_n(N) - T_n^*(N) > \varepsilon \right) = 0.$$

Since for all N , we have

$$\begin{aligned} \hat{T}_n(N) &\equiv \max_{i=1, \dots, n} \tilde{T}_n(N) \text{ and} \\ T_n^*(N) &\geq \max_{i=1, \dots, n} T_i^*(N), \end{aligned}$$

we have also

$$\lim_{N \rightarrow \infty} \Pr \left(\hat{T}_n(N) - T_n^*(N) > \varepsilon \right) = 0.$$

To show that $\Pr \left(T_n^*(N) - \hat{T}_n(N) > \varepsilon \right) \rightarrow 0$ is conceptually similar. ■

5 Simulating the informative equilibrium path

In this section we illustrate the main properties of the game by Monte-Carlo simulations. Proposition 4 gives a simple way to simulate the informative equilibrium directly in the limit $N \rightarrow \infty$. A sample path of the equilibrium is generated as follows. i) First, fix prior $\pi^0(\omega)$ and the true state of world ω' . ii) Draw a sequence $\{y_i\}_{i=1}^M$ of independent exponentially distributed random variables with parameter $g(\underline{\theta} | \omega')$. For this sequence, the corresponding sequence of posteriors is:

$$\pi^i(\omega) = \frac{\pi^{i-1}(\omega) g(\underline{\theta} | \omega) e^{-g(\underline{\theta} | \omega) \cdot y_i}}{\sum_{\omega \in \Omega} \pi^{i-1}(\omega) g(\underline{\theta} | \omega) e^{-g(\underline{\theta} | \omega) \cdot y_i}}, \quad i = 1, \dots, M.$$

For each $i = 1, \dots, M$, calculate the stopping time \hat{T}_i as:

$$\hat{T}_i = \max \left\{ \hat{T}_{i-1}, \arg \max_t \mathbb{E} \left[\sum_{\omega \in \Omega} \pi^i(\omega) \cdot v(t, \omega) \right] \right\}.$$

The generated sequence $\{\hat{T}_i\}_{i=1}^M$ is the simulated realization of the sequence (7) in the limit $N \rightarrow \infty$. By Proposition 4, it corresponds to the real time moments at which the first M players stop the game in the large game limit. By choosing M sufficiently large, one can ensure that the belief $\pi^M(\omega)$ has converged to the true state, i.e. $\pi^M(\omega') \approx 1$ and $\pi^M(\omega) \approx 0$ for all $\omega \neq \omega'$. This means that all the remaining players will stop in equilibrium at the same real time as the M :th player (with high probability). Thus, \hat{T}_M gives the real time at which the game collapses for this particular sample.

We illustrate next the model by Monte-Carlo simulations, where we generate a large number of equilibrium paths and use those to compute the probability distributions for the players' stopping times and payoffs. We specify the model as follows:

$$\begin{aligned}\Omega &= \left\{0, \frac{1}{S-1}, \frac{2}{S-1}, \dots, \frac{S-2}{S-1}, 1\right\}, \\ [\underline{\theta}, \bar{\theta}] &= [0, 1], \\ g(\theta|\omega) &= 1 + \gamma \left(\omega - \frac{1}{2}\right) \left(\theta - \frac{1}{2}\right), \\ v(t, \omega) &= -(\omega - t)^2.\end{aligned}$$

Here S is the number of states and γ is a parameter measuring the precision of individual signals. In this illustration we have $S = 10$, and for the signal precision we compare two cases: $\gamma = 2$ (precise signals) and $\gamma = 0.2$ (imprecise signals).

5.1 Distribution of stopping times

We generated 10000 sample paths for each 10 state values. For each sample path, we use $M = 300000$ random variables to make sure that the posteriors have fully converged to the true state. Figure 1 shows the simulated cumulative distribution functions of the moment of collapse, conditional on state. Top panel uses precision parameter $\gamma = 2$ while the bottom panel uses $\gamma = 0.2$. This Figure demonstrates clearly the Theorem 2: the time of collapse is random and delayed as compared to the first best for all but the highest state. The delay is more severe for the lowest state values. The signal precision has an expected effect: with less precise signals there is on average more delay.

Figure 1

5.2 Payoffs

Using the distributions of stopping times generated by the Monte-Carlo simulation, we can easily compute the ex-ante value of a player in equilibrium according to (3).

The following table shows the ex-ante values defined in section 4.2 and computed with the two precision parameters used in the simulations:

	V^0	V^1	V^*	V^P	V^S
$\gamma = 2$	-0.1019	-0.0984	-0.0690	0.0035	0.0294
$\gamma = 0.2$	-0.1019	-0.1018	-0.0989	0.000035	0.0029

The obvious result in this table is that the more precise the private signals, the more valuable private learning: V^P is higher for the precise signals. What is less obvious is

that the social value behaves similarly: the more precise the private signals, the more valuable is the additional value of the social learning on top of the private value of the signals. In fact, it is easy to show formally that in the limit where private signals are made uninformative in the sense that V^P goes to zero (in our model specification this would mean $\gamma \rightarrow 0$), also V^S must go to zero.

Figures 2 and 3 show the values conditional on signal and state, respectively. The value of an isolated player conditional on signal is U-shaped: extreme signal realizations are ex-ante good news in the quadratic payoff case, since they make large mistakes unlikely. In equilibrium, high signals are good news: they indicate that the optimal timing is more likely to be late, and social learning is particularly valuable if that is the case. Learning from others causes delay, which is valuable if late action is ex-post optimal, but it is costly if the early action would have been optimal. This can be seen more clearly in Figure 3 that shows the value functions conditional on state. Social learning makes payoffs more sensitive on true state: actions are delayed which is good if state is high but bad if state is low.

Figure 2

Figure 3

6 Discussion

Our results are quite different from related models in Chamley & Gale (1994) and Chamley (2004). To understand why this is the case, it is useful to note that we can embed the main features of those models as a special case of our model. For this purpose, assume that $\omega \in \{0, \infty\}$, and

$$v(t, 0) = e^{-rt}, v(t, \infty) = -ce^{-rt}.$$

If it is optimal to invest at all in this version of the model, then the investment time is insensitive to the information of the players. In other words, investment is good either immediately or never. Private signals only affect the relative likelihood of these two cases. This leads to the conclusion that it is never optimal to invest at $t > 0$ conditional on no other investments within $(t - \varepsilon, t)$, since then it would have been optimal to invest immediately. As a result, a given stage k ends either immediately if at least one player stops at time $t = 0$ and the play moves to stage $k + 1$, or the stage continues forever and the game never moves to stage $k + 1$. This means that all investment must take place at the beginning of the game, and with a positive probability investment stops forever even when $\omega = 0$. The models in Chamley & Gale (1994) and Chamley (2004) are formulated

in discrete time, but their limiting properties as the period length is reduced corresponds exactly to this description.

We get an intermediate case by setting $\Omega = \{\omega_1, \dots, \omega_S, \infty\}$ with $P(\infty) > 0$. In this case, the game has some revelation of information throughout the game. Nevertheless, it is possible that all investment ends even though $\omega < \infty$, and as a result, the game allows for a similar possibility of incorrect actions as Chamley & Gale (1994).


There are a number of directions where the analysis in this paper should be extended. Exogenous uncertainty on the payoff of investment plays an important role in the literature on real options. Our paper can be easily extended to cover the case where the profitability of the investment depends on an exogenous (and stochastic) state variable p and on private information about common market state ω . In this formulation, the stage game is one where the players pick a Markovian strategy for optimal stopping. With our monotonicity assumptions this is equivalent to selecting a threshold value $p_i(\theta_i)$ at which to stop conditional on their signal. The stage ends at the first moment when the threshold value of some player is hit.

The analytical simplicity of the model also makes it worthwhile to consider some alternative formulations. First, it could be that the optimal time to stop for an individual player i depends on the common parameter ω as well as her own signal θ_i . The reason for considering this extension would be to demonstrate that the form of information aggregation demonstrated in this paper is not sensitive to the assumption of pure common values. Second, by including the possibility of payoff externalities in the game we can bring the current paper closer to the auction literature. We plan to investigate these questions in future work.


References

- Bulow, J. & P. Klemperer. 1994. "Rational Frenzies and Crashes." *Journal of Political Economy* 102:1–23.
- Chamley, C. 2004. "Delays and Equilibria with Large and Small Information in Social Learning." *European Economic Review* 48:477–501.
- Chamley, C. & D. Gale. 1994. "Information Revelation and Strategic Delay in a Model of Investment." *Econometrica* 62:1065–1086.
- Feddersen, T. & W. Pesendorfer. 1997. "Voting Behavior and Information Aggregation in Elections With Private Information." *Econometrica* 65:1029–1058.
- Kremer, I. 2002. "Information Aggregation in Common Value Auctions." *Econometrica* 70:1675–1682.
- Levin, D. & J. Peck. 2008. "Investment Dynamics with Common and Private Values."
- Moscarini, G. & F. Squintani. 2008. "Competitive Experimentation with Private Information."
- Murto, P. & J. Välimäki. 2009. "Learning and Information Aggregation in an Exit Game."
- Pesendorfer, W. & J. Swinkels. 1997. "The Loser's Curse and Information Aggregation in Common Value Auctions." *Econometrica* 65(6):1247–1282.
- Smith, L. & P. Sorensen. 2000. "Pathological Outcomes of Observational Learning." *Econometrica* 68:371–398.


CDF of equilibrium stopping times for different state realizations. Precise signals.



CDF of equilibrium stopping times for different state realizations. Imprecise signals.



Expected payoff conditional on signal (Blue=equilibrium, Red=isolated player, Solid = precise signals, Dashed = imprecise signals)



Expected payoff conditional on state (Blue=equilibrium, Red=isolated player, Solid = precise signals, Dashed = imprecise signals)

