

Delay and Information Aggregation in Stopping Games with Private Information

Pauli Murto* Juuso Välimäki †

February 2011

Abstract

We consider equilibrium timing decisions in a model with informational externalities. A number of players have private information about a common payoff parameter that determines the optimal time to invest. The players learn from each other in a continuous-time multi-stage game by observing the past investment decisions. 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.

1 Introduction

This paper analyzes a game of timing where the players are privately informed about a common payoff parameter that determines the optimal time to stop the game. Information is transmitted across the players through observed actions, i.e. realized individual stopping decisions. In other words, our model

*Department of Economics, Aalto University, and HECER pauli.murto@aalto.fi.

†Department of Economics, Aalto University, and HECER juuso.valimaki@aalto.fi.

is one of observational learning, where communication between the players is not considered. For concreteness, we sometimes interpret the stopping decision as an irreversible investment decision as in the literature on real options.

The key question in our paper is how the individual players balance the benefits from observing other players' actions with the costs of delay. Observational learning is potentially socially valuable because it allows information to spread across players. When timing their own decisions, however, the players disregard the informational benefits to the other players. This informational externality leads to delays when contrasted with socially efficient information transmission. As a result, much of the potential value of social learning is dissipated.

Our main findings are: i) The most informative symmetric equilibrium results in delays. ii) The most informative symmetric equilibrium displays herding in the sense that when the number of players is large, almost all players stop at the same time. iii) Even in large games with accurate pooled information, aggregate uncertainty persists.

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 the players have information on ω , the 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 ratio property. The payoffs are assumed to be either supermodular or logsupermodular in ω and the stopping time t .

We show that the game has symmetric equilibria in monotone strategies. Our main characterization result describes a simple method for calculating the optimal stopping moment for each player in the most informative symmetric equilibrium of the game. In this equilibrium, a player always calculates her stopping time as if her own signal were the most extreme (that

is, favoring early investment) amongst those players that have not yet invested. The game has also less informative equilibria where all the players stop immediately regardless of their signals.

We allow the players to react quickly to each other's stopping decisions. In order to avoid complicated limiting procedures, we model the dynamic game as a multi-stage game with continuous action spaces. At the beginning of each stage, all the remaining players choose their stopping time from the real line. The stage ends at the minimum of these stopping times. This minimum stopping time and the identity of the player(s) that chose it are publicly observed. 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. Immediate reactions to stopping decisions are captured by allowing stopping time 0 in the next stage. We believe that this formulation can be useful in some other dynamic settings as well to simplify the definition of admissible strategies.

To understand the source of delay in our model, it is useful to point out an inherent asymmetry in learning in stopping games. While the players can decide at time t to delay their actions in response to new information, they cannot decide to go backward in time and stop at $t' < t$. Hence, it is possible to learn that stopping is taking place too late. By contrast, since waiting is always an option, it is not possible to become convinced that stopping is taking place too early.

We obtain the sharpest results for games with a large number of players. First, 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). Second, we show that almost all the players stop at the same instant of real time (even though they may stop in different stages) where the game also ends. Finally, we show that even if we condition on the true state, the time at which the

players stop remains stochastic.

To understand how the model works, consider the simplest case where the players observe a binary signal on the true state of the world. If player i is the only player in the game, she simply stops at the optimal moment given her posterior. Given that the signals satisfy MLRP and the payoff is supermodular in ω and t , then she stops earlier, say at t_L if her signal is low. Suppose next that there are $N > 1$ players and consider the incentives of the players that have received a low signal. If the other players with a low signal were to stop at t_L , then it would be in the best interest of player i to wait a bit longer to observe the decisions at t_L and hence to find out the number of low signals amongst the other players. This rules out an equilibrium where all players with low signals stop at t_L . On the other hand, it is also impossible that in a symmetric equilibrium no player stops with a positive probability at t_L . If this were the case, then the first player to stop after t_L would act upon the information contained in her own signal only. But with this information, the optimal stopping time is t_L . Hence in equilibrium, the benefits from learning from others must be balanced with the costs of delay.

Our main results characterize the speed of observational learning that achieves this balance. When the number of players is large, we show that the statistical informativeness of observational learning depends only on the conditional densities of the lowest possible signals. Because of the informational externalities, players stop too late relative to the full information benchmark. At the same time, most players gain significantly relative to stopping in isolation from the other players.

RELATED LITERATURE

Our paper is related to the literature on herding. Early papers such as Banerjee (1992) and Bikhchandani, Hirshleifer & Welch (1992) assumed an exogenous order of moves for the players. Like us, Grenadier (1999) relaxes this assumption in order to address observational learning in a model of investment. However, in his model players are exogenously ranked in terms

of the informativeness of their signals, and this ranking is common knowledge. This assumption plays a role similar to the assumption of exogenous order of moves, and as a result, the model features information cascades through a mechanism similar to Banerjee (1992) and Bikhchandani, Hirshleifer & Welch (1992). By contrast, we assume that the players are ex-ante similar, and this leads to qualitatively different pattern of information revelation. Our model has no information cascades, but information is revealed inefficiently late.

The paper that is closest to ours is the model of entry by Chamley & Gale (1994).¹ The key difference to our paper is that in their model it is optimal to invest either 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. In other words, Chamley & Gale (1994) models uncertainty over *whether* it is optimal to invest or not, while we model uncertainty over *when* it is optimal to invest. This turns out to have important implications for the model's predictions. 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.² The information revelation in our model is closely related to our previous paper Murto & Välimäki (2011). In that paper, private learning over time generates dispersed information about the optimal stopping point, and information aggregates in sudden bursts of action. Moscarini & Squintani (2010) analyze a two-firm R&D race where the inference on common values information is similar to our model. The results and the analysis in the two papers are quite different since our main focus is on information aggregation in a general class of stopping models with pure informational externalities.

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

¹See also Chamley (2004) for a more general model. Levin & Peck (2008) extends this type of a model to allow private information on the stopping cost. In contrast to our model, information is of the private values type in their model.

²Section 6 discusses in more detail the relationship between these papers.

with common values, Pesendorfer & Swinkels (1997) show that information aggregates efficiently as the number of objects 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 a failure of information aggregation even for large economies. On the other hand, Bulow & Klemperer (1994) analyzes an auction model that features "frenzies" that resemble equilibrium stopping behavior in our model. In Bulow & Klemperer (1994) those are generated by direct payoff externalities arising from scarcity, whereas our equilibrium dynamics relies on a purely informational mechanism.

The paper is structured as follows. Section 2 introduces the basic model. Section 3 establishes the existence of a symmetric monotonic equilibrium. Section 4 discusses the properties of the game with a large number of players. Section 5 presents a quadratic example of the model. Section 6 compares our results to the most closely related literature. Section 7 concludes.

2 Model

2.1 Payoffs and signals

N players consider investing in a project. The payoff for player i from an investment at time t_i depends on the state $\omega \in \Omega$, and is given by a continuous function

$$v(t_i, \omega).$$

The state space is a compact set $\Omega \subseteq [0, 1]$, and can be either finite or infinite.³ The players share a common prior $p^0(\omega)$ on Ω . The players choose

³We can transform the state space to the extended real line by a monotone transformation $x = \frac{\omega}{1-\omega}$. With this transformation, the state $\omega = 1$ corresponds to the case where it is never optimal to stop.

their investment time t from the set $T = [0, \infty)$.

The players face uncertainty over ω and choose the timing of their investment in order to maximize their expectation of v . We assume that the payoff function v is twice differentiable in t almost everywhere and either supermodular or log-supermodular in (t_i, ω) :

Assumption 1 *Either,*

$$v(t_i, \omega) + v(t'_i, \omega') > v(t'_i, \omega) + v(t_i, \omega') \text{ for all } t_i > t'_i \text{ and all } \omega > \omega',$$

or

$$v(t_i, \omega)v(t'_i, \omega') > v(t'_i, \omega)v(t_i, \omega') \text{ for all } t_i > t'_i \text{ and all } \omega > \omega'.$$

This assumption is satisfied by all concave functions and all exponentially discounted concave functions.

We also assume that for each ω , there is a unique t that maximizes $v(t, \omega)$. By our assumption of strict (log-)supermodularity, this maximizer must be strictly increasing in ω . Therefore, without further loss of generality, we associate the state value directly with the first-best optimal investment time. In other words, we assume that v takes a form where the optimal stopping time for each state is equal to the state:

$$\omega = \arg \max_t v(t, \omega).$$

The players are initially privately informed about ω . Player i observes a signal $\theta_i \in \Theta = [0, \bar{\theta})$ for some $\bar{\theta} \leq \infty$. $G(\theta, \omega)$ is the joint probability distribution on $\Theta \times \Omega$. 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 have full support for all ω . We also assume that for all ω , $G(\theta | \omega)$ is continuous (i.e., there are no mass points) and $g(\theta | \omega)$ has at most a finite number of points of discontinuity and is continuous at $\theta = 0$.

We assume that the signals in the support of the signal distribution satisfy monotone likelihood property (MLRP):

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

$$\frac{g(\theta' | \omega')}{g(\theta | \omega')} \geq \frac{g(\theta' | \omega)}{g(\theta | \omega)}. \quad (1)$$

This assumption, together with Assumption 1, allows us to conclude that the optimal stopping time conditional on a signal is monotonic in the signal realization. That is, player i 's optimal stopping time is increasing in her own type as well as in the type of any other player j .

Finally, we make an assumption for the signal densities at the lower end of the signal distribution. This assumption has two purposes. First, we want to make sure that the signals can distinguish different states. This is guaranteed by requiring $g(0|\omega) \neq g(0|\omega')$ whenever $\omega \neq \omega'$ (note that assumption 2 alone allows conditional signal densities that are identical in two different states). Second, we want to rule out the case where some players can infer the true state from observing their own signal. This is guaranteed by requiring $g(0|\omega) < \infty$ for all $\omega \in \Omega$. While none of the players can infer the true state based on their own signal, the assumption of conditionally independent signals and MLRP together guarantee that the pooled information held by the players becomes arbitrarily informative as the number of players tends to infinity.

Assumption 3 For all $\omega, \omega' \in \Omega$, $\omega' > \omega$,

$$0 < g(0|\omega') < g(0|\omega) < \infty.$$

2.2 Strategies and information

We assume that at t , the players know their own signals and the past decisions of the other players. We do not want our results to depend on any exogenously set observation lag. Therefore, we allow the players to react immediately to new information that they obtain by observing that other players stop the game. To deal with this issue in the simplest manner, we model the game as a multi-stage stopping game as follows.

The game consists of a random number of stages with partially observable actions. In stage 0, all players choose their investment time $\tau_i(h^0, \theta_i) \in [0, \infty)$ depending on their signal θ_i . The stage ends at $t^0 = \min_i \tau_i(h^0, \theta_i)$. At that point, the set of players that invest at t^0 , i.e. $\mathcal{S}^0 = \{i : \tau_i(h^0, \theta_i) = t^0\}$ is announced. The actions of the other players are not observed. The public history after stage 0 and at the beginning of stage 1 is then $h^1 = (t^0, \mathcal{S}^0)$. The vector of signals θ and the stage game strategy profile $\tau(h^0, \theta) = (\tau_1(h^0, \theta_1), \dots, \tau_N(h^0, \theta_N))$ induce a probability distribution on the set of histories H^1 . The public posterior on Ω (conditional on the public history only) at the end of stage 0 is given by Bayes' rule:

$$p^1(\omega | h^1) = \frac{p^0(\omega) \Pr(h^1 | \omega)}{\int_{\Omega} p^0(\omega) \Pr(h^1 | \omega) d\omega}.$$

As soon as stage 0 ends, the game moves to stage 1, which is identical to stage 0 except that the set of active players excludes those players that have already stopped. Once stage 1 ends, the game moves to stage 2, and so forth. In general, in each stage k , players that have not yet invested choose an investment time $\tau_i(h^k, \theta_i) \geq 0$. We let \mathcal{N}^k denote the set of players that are still active at the beginning of stage k (i.e., players that have not yet stopped in stages $k' < k$). The public history available to the players is

$$h^k = h^{k-1} \cup (t^{k-1}, \mathcal{S}^{k-1}).$$

The set of stage k histories is denoted by H^k , and the set of all histories by $H := \cup_k H^k$. We denote the number of players that invest in stage k by S^k and the cumulative number of players that have invested in stage k or earlier by $Q^k := \sum_{i=0}^k S^i$.

A pure strategy in the game is a function

$$\tau_i : H \times \mathbb{R}_+ \rightarrow [0, \infty).$$

The duration of stage k is given by

$$t^k = \min_{i \in \mathcal{N}^k} \tau_i(h^k, \theta_i).$$

Throughout the paper, we use the term *real time* to refer to the elapse of time from the beginning of the game. If player i is among those who stop in stage k , i.e. if $\tau_i(h^k, \theta_i) = t^k$, then the realized payoff of i is $v(T^k + t^k, \omega)$, where T^k records the real time at the beginning of stage k :

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

The players maximize their expected payoff. A strategy profile $\tau = (\tau_i, \dots, \tau_N)$ is a Perfect Bayesian Equilibrium of the game if for all i and all θ_i and h^k , $\tau_i(h^k, \theta_i)$ is a best response to τ_{-i} .

3 Monotonic Symmetric Equilibrium

In this section, we analyze symmetric equilibria in monotonic pure strategies.

Definition 1 *A strategy τ_i is monotonic if for all k and h^k , $\tau_i(h^k, \theta)$ is (weakly) increasing in θ .*

With a monotonic symmetric strategy profile, the players stop the game in the increasing order of their signal realizations. Therefore, at the beginning of stage k , it is common knowledge that all the remaining players have signals within $(\underline{\theta}^k, 1)$, where:

$$\underline{\theta}^k := \max \{ \theta \mid \tau(h^{k-1}, \theta) = t^{k-1} \}. \quad (2)$$

3.1 Informative Equilibrium

In this section, we characterize the symmetric equilibrium that maximizes information transmission in the set of symmetric monotone pure strategy equilibria. Theorem 1 states that there is a symmetric equilibrium, where a player with the signal θ stops at the optimal time conditional on all the other active players having a signal at least as high as θ . The monotonicity of

this strategy profile follows from MLRP. We call this profile the informative equilibrium of the game.

To state the result, we define the smallest signal among the active players at the beginning of stage k :

$$\theta_{\min}^k := \min_{i \in \mathcal{N}^k} \theta_i.$$

Theorem 1 (Informative equilibrium) *The game has a symmetric equilibrium profile τ^* in monotonic strategies, where the stopping time for a player with signal θ at stage k is given by:*

$$\tau^*(h^k, \theta) := \min \left(\arg \max_{t \geq 0} \mathbb{E} [v(t + T^k, \omega) | h^k, \theta_{\min}^k = \theta] \right). \quad (3)$$

The proof is in the appendix, and it utilizes the key properties of $\tau^*(h^k, \theta)$ stated in the following Proposition:

Proposition 1 (Properties of informative equilibrium) *The stopping time $\tau^*(h^k, \theta)$ defined in (3) is increasing in θ . Furthermore, for $k \geq 1$, there is some $\varepsilon > 0$ such that along equilibrium path, $\tau^*(h^k, \theta) = 0$ for all $\theta \in [\underline{\theta}^k, \underline{\theta}^k + \varepsilon)$.*

Proof. Proposition 1 is proved in the Appendix. ■

The equilibrium stopping strategy $\tau^*(h^k, \theta)$ defines a time-dependent cutoff signal $\theta^{*k}(t)$ for all $t \geq 0$:

$$\theta^{*k}(t) := \max \{ \theta | \tau^*(h^k, \theta) \leq t \}. \quad (4)$$

In words, $\theta^{*k}(t)$ is the highest type that stops at time t in equilibrium. Proposition 1 implies that along the informative equilibrium path, $\theta^k(0) > \underline{\theta}^k$ for all stages except possibly the first one. This means that all the players with a signal in the interval $(\underline{\theta}^k, \theta^k(0))$ stop immediately at the beginning of the stage, and there is therefore a strictly positive probability that many players stop simultaneously.

To understand the equilibrium dynamics in stage k , note that as real time moves forward, the cutoff $\theta^{*k}(t)$ truncates the interval within which the signals of the remaining players lie. By MLRP and the (log)supermodularity of v , this new information delays the optimal stopping time for all the remaining players. At the same time, the passage of time increases the relative payoff from stopping the game for each signal θ . In equilibrium, $\theta^{*k}(t)$ increases at a rate that balances these two effects and keeps the marginal type indifferent. As soon as the stage ends at $t^k > 0$, the players learn that one of the other active players has a lower than expected signal. By MLRP and the (log)supermodularity of v , the expected value from staying in the game falls by a discrete amount for the remaining players. This means that the marginal cutoff moves discretely upwards and explains why $\theta^{*(k+1)}(0) > \theta^k(t^k) = \underline{\theta}^{k+1}$. As a result, each new stage begins with a positive probability of immediate further exits. If at least one player stops so that $t^{k+1} = 0$, the game moves immediately to stage $k + 2$. The preceding argument can be repeated until there is a stage with no further immediate exits. Thus, the equilibrium path alternates between *stopping phases*, i.e. consecutive stages k' that end at $t^{k'} = 0$ and that result in multiple simultaneous exits, and *waiting phases* where all players stay in the game for time intervals of positive length.

Note that the random time at which stage k ends,

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

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 , $\tau^*(h^k, \theta_i)$ 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 realizations of all previous statistics is known. Unfortunately this is not the case since for all $k > 1$, there is a strictly positive mass of types that stop immediately at $t^k = 0$. This implies that the signals of those players that stop immediately are imperfectly revealed in equilibrium. However, in Section 4.1 we show that in the limit as the number of players

is increased towards infinity, payoff relevant information in equilibrium converges to the payoff relevant information contained in the order statistics of the signals.

3.2 Uninformative equilibria

Some stage games also have an additional symmetric equilibrium. In these equilibria, players use strategies that do not depend on their signals. We call these equilibria *uninformative*. They are similar to *rush* equilibria in Chamley (2004).

To understand when such uninformative equilibria exist, consider the optimal stopping problem of a player who conditions her decision on history h^k and her private signal θ_i , but not on the other players having signals higher than hers. If $t = 0$ solves that problem for all signal types remaining in the game, i.e., if

$$0 \in \arg \max_{t \geq 0} \mathbb{E} [v(t + T^k, \omega) | h^k, \theta_i = \theta] \text{ for all } \theta \geq \underline{\theta}^k,$$

then an uninformative equilibrium may exist. If all players stop at $t = 0$ then they learn nothing from each other. If they learn nothing from each other, then $t = 0$ is their optimal action.

It should be noted that some equilibria where all the players stop immediately satisfy our criteria for informative equilibrium. If $\tau^*(h^k, \theta) = 0$ for all θ , then the continuation equilibrium is informative in our terminology even though all players stop at once. At any such history h^k , the players find it optimal to exit even if all the remaining players had the highest possible signal. Similarly, with some payoff specifications there are informative equilibria where all the players stop at $t = 1$ even if their beliefs have not converged (which, in such a case, is to be interpreted as delaying infinitely). See discussion of such a case in Section 6.

In the *least informative equilibrium*, uninformative equilibrium is played in *all* stages where the above criterion is satisfied. There are also intermediate

equilibria where after some h^k , players use $\tau^*(h^k, \theta)$ defined in (3), and after other h^k , they play uninformatively.

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. This follows from the fact that every player can always ensure the outcome of the uninformative equilibrium after all h^k regardless of the other players' strategy choices.

4 Informative Equilibrium in Large Games

In this section, we study the limiting properties of the model as we increase the number of players towards infinity. We will see that those properties can be described by means of the order statistics of the players' signals. We start with a simple statistical observation regarding the distribution of order statistics in large samples.

4.1 Information in equilibrium

Denote the n^{th} order statistic in the game with N players by

$$\tilde{\theta}_n^N := \min \{ \theta \in [0, \bar{\theta}] \mid \# \{ i \in \mathcal{N} \mid \theta_i \leq \theta \} = n \}. \quad (5)$$

It is clear that if we increase N towards infinity while keeping n fixed, $\tilde{\theta}_n^N$ converges to the lower bound 0 of the signal distribution in probability. Therefore, we scale the order statistics by the number of players:

$$Z_n^N := \tilde{\theta}_n^N \cdot N. \quad (6)$$

Since Z_n^N is a deterministic function of $\tilde{\theta}_n^N$, it has the same information content as $\tilde{\theta}_n^N$. Nevertheless, it is a well known statistical fact that Z_n^N 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.

We define also

$$Y_n^N := Z_n^N - Z_{n-1}^N = \left(\tilde{\theta}_n^N - \tilde{\theta}_{n-1}^N \right) \cdot N, \quad (7)$$

where by convention we let $\theta_0^N = Z_0^N = 0$. Let $[Y_1^\infty, \dots, Y_n^\infty]$ be a vector of n independent exponentially distributed random variables with parameter $g(0 | \omega)$:

$$\Pr(Y_1^\infty \leq x_1, \dots, Y_n^\infty \leq x_n) = e^{-g(0|\omega) \cdot x_1} \cdot \dots \cdot e^{-g(0|\omega) \cdot x_n}.$$

Proposition 2 Fix $n \in \mathbb{N}_+$. Consider the sequence of random variables $\{[Y_1^N, \dots, Y_n^N]\}_{N=n}^\infty$, where for each N the random variables Y_i^N are defined by (5) - (7). As $N \rightarrow \infty$, we have:

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

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

Proof. In the Appendix. ■

Proposition 2 states that in the limit $N \rightarrow \infty$, learning from the order statistics is equivalent to sampling independent random variables from an exponential distribution with an unknown state-dependent parameter $g(0 | \omega)$. The intuition is straight-forward. As N is increased, the n lowest order statistics converge towards 0. Therefore, the signal densities matter for the learning only in the limit $\theta \downarrow 0$. One can think of $g(0 | \omega)$ as the intensity of the order statistics in the large game limit. This explains why we have adopted the assumption that the signal density $g(\theta | \omega)$ is continuous at $\theta = 0$.

Note that $Z_n^N = \sum_{i=1}^n Y_i^N$ converges to a sum of independent exponentially distributed random variables, which means that the limiting distribution of Z_n^N is Gamma distribution:

Corollary 1 For all n ,

$$Z_n^N := \sum_{i=1}^n Y_i^N \xrightarrow{\mathcal{D}} Z_n^\infty,$$

where $Z_n^\infty \sim \Gamma(n, g(0 | \omega))$.

By Proposition 2, observing the n lowest order statistics is equivalent to observing n independent exponentially distributed random variables when N is large. By the memoryless property of the exponential distribution, observing only the n^{th} order statistic $\tilde{\theta}_n^N$ is informationally equivalent to observing all order statistics up to n . To see this formally, denote by $\pi(\omega | (y_1, \dots, y_n))$ the posterior probability of an arbitrary element $\omega \in \Omega$ based on an independent sample $Y_1^\infty = y_1, \dots, Y_n^\infty = y_n$, and let $\pi(\omega | z_n)$ denote the corresponding posterior probability based on the sample that contains only the sum of the previous sample, $z_n = z_1 + \dots + z_n$. Bayesian rule and simple algebra show that these posteriors are equal:

$$\begin{aligned} \pi(\omega | (y_1, \dots, y_n)) &= \frac{\pi^0(\omega) \cdot \prod_{i=1}^n g(0 | \omega) e^{-g(0|\omega)(y_i - y_{i-1})}}{\int_{\Omega} \pi^0(\omega) \cdot \prod_{i=1}^n g(\theta | \omega) g(0 | \omega) e^{-g(0|\omega)(y_i - y_{i-1})} d\omega} \\ &= \int_{\Omega} \frac{\pi^0(\omega) \cdot (g(0 | \omega))^n e^{-g(0|\omega)z_n}}{\pi^0(\omega) \cdot (g(0 | \omega))^n e^{-g(0|\omega)z_n}} d\omega = \pi(\omega | z_n). \end{aligned}$$

In the finite model, the posterior $\pi^N(\omega | (y_1, \dots, y_n))$ based on a sample $Y_1^N = y_1, \dots, Y_n^N = y_n$ generally differs from the posterior $\pi^N(\omega | z_n)$ that is based only on $z_n = y_1 + \dots + y_n$. Nevertheless, the Bayes' rule is continuous in the limit as $N \rightarrow \infty$ in (y_1, \dots, y_n) since we assume $g(\cdot | \omega)$ to be continuous at $\theta = 0$ for all ω . Therefore, Proposition 2 implies that both $\pi^N(\omega | (y_1, \dots, y_n))$ and $\pi^N(\omega | z_n)$ converge to the posterior $\pi(\omega | z_n)$ for all ω and (y_1, \dots, y_n) as $N \rightarrow \infty$. We summarize this discussion in the following Corollary.

Corollary 2 *For a fixed sample of normalized order statistics (y_1, \dots, y_n) ,*

$$\lim_{N \rightarrow \infty} \pi^N(\omega | (y_1, \dots, y_n)) = \lim_{N \rightarrow \infty} \pi^N(\omega | z_n) = \pi(\omega | z_n) \text{ for all } \omega,$$

where $z_n = y_1 + \dots + y_n$.

This corollary plays a key role in our analysis. Suppose that player i has a signal θ and that she observes all the signals below θ . By Theorem 3,

she would choose the optimal stopping time conditional on the values of the lower signals and conditional on the assumption that all other players have signals above θ (and of course subject to the restriction that stopping before the current instant of real time is impossible). Corollary 2 implies that the number of players n with signals below θ summarizes the relevant part of the history in the limit as $N \rightarrow \infty$. Hence even if all signals were observable, the relevant conditioning event is $Z_n^N = N\theta$ when $N \rightarrow \infty$.

4.2 Timing in Large Games

So far, we have discussed the properties of the order statistics without linking them to the optimal timing decisions. In this section we consider the properties of the informative equilibrium, and show that when $N \rightarrow \infty$, the equilibrium path of the game can be approximated by a simple algorithm that samples sequentially the order statistics.

We consider first the hypothetical case, where a decision maker chooses the optimal timing based on the information contained in the order statistics of the limit model. Let $U(t | z_n)$ denote the expected utility from stopping at time t , given information that $Z_n = z_n$ with $Z_n \sim \Gamma(n, g(0 | \omega))$:

$$U(t | z_n) := \int_{\Omega} v(t, \omega) \pi(\omega | z_n) d\omega.$$

As a preliminary result, we prove the uniqueness of the optimal solution to this problem for almost every realization z_n of Z_n .

Lemma 1 *Let $Z_n \sim \Gamma(n, g(0 | \omega))$ and define*

$$t_n(z_n) := \arg \max U(t | z_n).$$

Then $t_n(z_n)$ is a singleton for almost every z_n in the measure induced by the random variable Z_n on \mathbb{R}_+ .

Proof. In the Appendix. ■

We turn next to the finite model with N players. Consider a sample of normalized order statistics

$$(Y_1^N = y_1, \dots, Y_n^N = y_n),$$

and let $U^N(t \mid (y_1, \dots, y_n))$ and $U^N(t \mid z_n)$ denote the expected utilities from stopping at time t , based on the whole sample (y_1, \dots, y_n) and sample $z_n = y_1 + \dots + y_n$, respectively:

$$U^N(t \mid (y_1, \dots, y_n)) := \int_{\Omega} v(t, \omega) \pi^N(\omega \mid (y_1, \dots, y_n)) d\omega,$$

$$U^N(t \mid z_n) := \int_{\Omega} v(t, \omega) \pi^N(\omega \mid z_n) d\omega.$$

The corresponding optimal stopping time are:

$$t_n^N(y_1, \dots, y_n) : = \arg \max U^N(t \mid (y_1, \dots, y_n)),$$

$$t_n^N(z_n) : = \arg \max U^N(t \mid z_n).$$

Note that $t_n^N(y_1, \dots, y_n)$ and $t_n^N(z_n)$ could in principle be sets. However, the next lemma says that they both converge to $t_n(z_n)$, which is singleton for almost every z_n by Lemma 1.⁴

Lemma 2 *For almost every (y_1, \dots, y_n) , $z_n = y_1 + \dots + y_n$,*

$$\lim_{N \rightarrow \infty} t_n^N(y_1, \dots, y_n) = \lim_{N \rightarrow \infty} t_n^N(z_n) = t_n(z_n).$$

Proof. In the Appendix. ■

With this Lemma, we can prove that the equilibrium stopping times converge to the stopping times of the limit model for almost all z_n . Notice that the unconstrained optimal times $t_k(z_k)$ could be decreasing in k . Since the players cannot go backwards in time, the relevant (constrained) stopping

⁴Clearly, since $t_n^N(y_1, \dots, y_n)$ and $t_n^N(z_n)$ converge to the same value for almost every (y_1, \dots, y_n) , no partial information on y_1, \dots, y_n could change the optimal timing decision of a Bayesian decision maker who has observed z_n (in the limit $N \rightarrow \infty$).

time in the limit model for the player with n^{th} lowest signal is the maximum of $t_{n'}(z_{n'})$, $n' = 1, \dots, n$:

$$T_n(y_1, \dots, y_n) := \max_{n'=1, \dots, n} t_{n'}(z_{n'}). \quad (8)$$

The main result of this section is that the stopping times in the informative equilibrium of the game converge to the stopping times defined in (8). We denote by $T_n^N(y_1, \dots, y_n)$ the real time at which the player with the n^{th} lowest signal stops in the informative equilibrium.

Proposition 3 *For all n , and for almost every (y_1, \dots, y_n) ,*

$$\lim_{N \rightarrow \infty} T_n^N(y_1, \dots, y_n) = T_n(y_1, \dots, y_n).$$

Proof. In the Appendix. ■

We end this section with a corollary that relates the distribution of realized stopping times in the symmetric equilibrium to the realized stopping times in the limit model. Let T_n^N denote the random stopping time of the n^{th} player in the symmetric equilibrium and let T_n denote the optimal stopping time based on the n^{th} lowest order statistic in the limit model as defined in (8).

Corollary 3 *The realized stopping times in the symmetric equilibrium converge in distribution to the stopping times in the limit model:*

$$T_n^N(Y_1^N, \dots, Y_n^N) \xrightarrow{\mathcal{D}} T_n(Y_1^\infty, \dots, Y_n^\infty).$$

Proof. Direct consequence of Propositions 2 and 3. ■

4.3 Delay in Equilibrium

In this section, we characterize the real time behavior of the players in the informative equilibrium in the limit as $N \rightarrow \infty$. 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 beginning of the game is N . We will be particularly interested in the behavior of $T_N(\theta, \omega)$ as N grows and we define

$$T(\omega, \theta) := \lim_{N \rightarrow \infty} T^N(\omega, \theta),$$

where the convergence is to be understood in the sense of convergence in distribution.

The real time instant at which the last player invests is denoted by $T^N(\omega)$ and we let

$$T(\omega) := \lim_{N \rightarrow \infty} T^N(\omega).$$

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. By $t(0)$, we denote the optimal investment time of a player that decides based on signal $\theta = 0$ only.

Theorem 2 *In the informative equilibrium of the game, we have for all $\omega < \max \Omega$,*

1. *supp $f(t | \omega) = [\max\{t(0), \omega\}, \max \Omega]$.*
2. *For all $\theta, \theta' > 0$,*

$$\lim_{N \rightarrow \infty} \Pr\{|T^N(\omega, \theta) - T^N(\omega, \theta')| < \varepsilon\} = 1 \text{ for all } \varepsilon > 0.$$

Proof. In the Appendix. ■

Theorem 2 confirms the main properties of our model. Almost all the players stop (almost) simultaneously (by Part 2 of the theorem), and this

stopping moment is inefficiently late and random (by Part 1 of the theorem). Since all the players with signals strictly above zero stop at the same time, the informational properties of the model are driven by the lowest signals. All the relevant information is transmitted by the lowest order statistics, and it is irrelevant how good information might be available at higher signal values.

5 Example with quadratic payoffs

In this section, we compute analytically the statistical properties of the informative equilibrium in the large-game limit for a special case of our model. As in much of the literature on observational learning, we assume that both the states and the signals are essentially binary. There are N ex ante identical players. We let $\omega \in \{0, 1\}$ and we map the binary signal setting into our model by assuming the following signal densities:

$$\frac{g(\theta | 0)}{g(\theta | 1)} = c_l \quad \text{for all } 0 \leq \theta \leq \theta^*, \quad (9)$$

$$\frac{g(\theta | 0)}{g(\theta | 1)} = c_h \quad \text{for all } \theta^* < \theta < \bar{\theta}. \quad (10)$$

Hence all the signals below (above) θ^* have the same informational content defined by parameter c_l (c_h). Sometimes we call signals below (above) θ^* low (high) and write $\theta = l(= h)$. For simplicity, we assume that

$$G(\theta^*, 0) = 1 - G(\theta^*, 1) = \alpha > \frac{1}{2},$$

which implies that $c_l = \alpha / (1 - \alpha)$ and $c_h = (1 - \alpha) / \alpha$. Hence, α measures the precision of the signals. We also assume that the prior probability $p^0 = \Pr\{\omega = 1\} = \frac{1}{2}$.

The payoffs in the model are given by

$$u(t, \omega) = -(t - \omega)^2.$$

Hence the optimal action for a player with posterior p on $\{\omega = 1\}$ is to invest at $t = p$.

We start the analysis by calculating the payoffs of a player that decides the timing of her investment in isolation of other players. At the prior stage, her payoff is

$$V^0 = \frac{1}{2} \left(-\frac{1}{4} \right) + \frac{1}{2} \left(-\frac{1}{4} \right) = -\frac{1}{4}.$$

After observing her signal, her posterior becomes more informative. If she observes a signal $\theta \leq \theta^*$, her posterior becomes $p = 1 - \alpha$. If $\theta > \theta^*$, her posterior is $p = \alpha$. Hence her payoff after observing her own signal is

$$\begin{aligned} V^I &= -\alpha(1 - \alpha)^2 - (1 - \alpha)\alpha^2 \\ &= -\alpha(1 - \alpha). \end{aligned}$$

Notice that the loss from the decisions vanishes as the signals get accurate, i.e. $V^I \uparrow 0$ as $\alpha \uparrow 1$. On the other hand, as $\alpha \downarrow 1/2$, signals become uninformative and $V^I \downarrow V^0 = -1/4$.

Consider next the case with a large N . If all the players were able to pool their information, then the posterior would be very informative of the true state, and all the players would stop together at the efficient stopping time. This follows from the fact that the number of players with a signal below θ^* is a binomial random variable $X^0(N)$ (or $X^1(N)$) with parameter α (or $1 - \alpha$) if $\omega = 0$ (or $\omega = 1$). We next investigate how well the players do if they can only observe each others' investment decisions but not their signals. That is, we consider the payoffs of the players in the informative equilibrium of the game.

From Theorem 3, we know that there is an informative equilibrium that is symmetric and in monotonic pure strategies. We denote this strategy profile by τ^* and the corresponding ex-ante payoff by V^* (this is the expected equilibrium payoff *prior* to observing the private signal θ).

When a player with signal θ^* invests, she behaves at every stage as if she knew that all other players have signals (strictly) above θ^* with probability 1 (again, this follows from Theorem 3). In order to compute V^* , we compute first the payoff of a player with signal θ that deviates to the strategy $\tilde{\tau} = \tau^*(h, \theta^*)$ for all $h \in H$. In other words, the deviating player just follows the strategy of the highest possible low signal player. We denote the ex ante expected payoff to the deviating player by \tilde{V} when all other players use their equilibrium strategies. Clearly this gives us a lower bound for V^* .

Denote by \tilde{T} the random real time at which the deviating player invests when using strategy $\tilde{\tau}$. Suppose that $\omega = 1$. Then $\omega = \max \Omega$, and Part 1 of Theorem 2 states that in the large game limit the last player stops at time $t = 1$. Part 2 of the same Theorem says that the stopping times of all signal types converge in probability to the same real time, hence we must have $\tilde{T} \rightarrow 1$ in probability. Therefore, denoting the expected payoff conditional on state ω by V_ω , we have:

$$\tilde{V}_1 \rightarrow 0$$

(in probability) as $N \rightarrow \infty$.

We turn next to the computation of \tilde{V}_0 . To do this, we define first the expected payoff $\tilde{V}_{\theta=l}$ of the deviating player when her signal is low, i.e. when $\theta < \theta^*$. Since the informational content of each such signal is the same and since the signals across players are conditionally independent, we know that this expected payoff is the same as the payoff to the player with the lowest possible signal $\theta = 0$. Since the player with the lowest signal is the first to invest in the informative equilibrium, her payoff is the same as the payoff based on her own signal only, and thus

$$\tilde{V}_{\theta=l} = V^I = -\alpha(1 - \alpha). \tag{11}$$

On the other hand, the probability of state $\omega = 0$ conditional on a low signal is α , and therefore

$$\tilde{V}_{\theta=l} = \alpha \tilde{V}_0 + (1 - \alpha) \tilde{V}_1. \quad (12)$$

Combining (11) and (12), and solving for \tilde{V}_0 gives:

$$\tilde{V}_0 = - (1 - \alpha) \left(1 + \frac{\tilde{V}_1}{\alpha} \right).$$

Therefore,

$$\begin{aligned} \tilde{V} &= \frac{1}{2} \tilde{V}_0 + \frac{1}{2} \tilde{V}_1 \\ &= -\frac{1 - \alpha}{2} + \left(\frac{1}{2} - \frac{1 - \alpha}{\alpha} \right) \tilde{V}_1 \\ &\rightarrow -\frac{1 - \alpha}{2} \text{ as } N \rightarrow \infty. \end{aligned}$$

The final step is to observe that as $N \rightarrow \infty$, we have $\tilde{V} \rightarrow V^*$ in probability. This follows from part 2 of Theorem 2, which says that the real stopping time of all signal types (except zero-probability case $\theta = 0$) converge to the same time, and therefore, the deviation that we have considered will not affect the realized payoff in the large game limit. Therefore, as $N \rightarrow \infty$,

$$V^* \rightarrow -\frac{1 - \alpha}{2}.$$

Note that we have $0 > V^* > V^I$ whenever $\alpha \in (\frac{1}{2}, 1)$. This means that the players benefit from the observational learning in equilibrium ($V^* > V^I$), but their payoff is nevertheless below efficient information sharing benchmark due to the informational externality ($V^* < 0$). Furthermore, it should be noted that $V_1^* \rightarrow 0$ and $V_0^* \rightarrow -(1 - \alpha)$. That is, observational learning benefits the players when $\omega = 1$, but hurts them when $\omega = 0$. Figure 1 draws the payoffs as functions of α .

< Figure 1 here >

To complete the analysis of the quadratic case, we analyze the distribution of \tilde{T} . As long as $t > 1 - \alpha$, but some of the uninformed players stay in the

game, they must be indifferent between staying and investing. Therefore, we must have

$$p_{\theta=l}(t) = t \text{ for all } t > 1 - \alpha,$$

where $p_{\theta=l}(t)$ denotes probability that a player with a low signal assigns on the event $\{\omega = 1\}$ at real time t . We already concluded that $\tilde{T} \rightarrow 1$ in probability if $\omega = 1$, and therefore, if it turns out that $\tilde{T} < 1$, then know that $p_{\theta=l}(t) = 0$ for all $t > \tilde{T}$. Therefore, we can compute the hazard rate $\chi_{\tilde{T}}(t)$ for the investment of the last player with a low signal in the limit as $N \rightarrow \infty$ from the martingale property of beliefs:

$$t = p_{\theta=l}(t) = (1 - \chi_{\tilde{T}}(t)dt)p_{\theta=l}(t + dt) + \chi_{\tilde{T}}dt \cdot 0,$$

or

$$\chi_{\tilde{T}}(t) = \frac{1}{t}.$$

Since $\Pr\{\tilde{T} < 1 - \epsilon | \omega = 1\} \rightarrow 0$ as $N \rightarrow \infty$, we can write the conditional probabilities of the event $\{\tilde{T} \in [t, t + dt) | \tilde{T} \geq t\}$ as

$$\begin{aligned} \chi_{\tilde{T}}(t \mid \omega = 0) &= \begin{cases} 0 & \text{for } t < 1 - \alpha \\ \frac{1}{t(1-t)} & \text{for } 1 - \alpha \leq t < 1 \end{cases} \\ \chi_{\tilde{T}}(t \mid \omega = 1) &= 0 \text{ for } t < 1. \end{aligned}$$

By Theorem 2, the probability distribution that we have derived for \tilde{T} is also the probability distribution for the stopping time of the last player in the game, which we have denoted $F(t \mid \omega)$. Figure 2 draws $F(t \mid 0)$ with different values of α .

< Figure 2 here >

To summarize, this quadratic binary example has demonstrated the following properties of our model: i) Observational learning is beneficial in high states and harmful in low states. ii) Inefficient delays persist for all but the highest state. iii) Almost all players invest at the same time as $N \rightarrow \infty$.

iv) The instant at which almost all players invest arrives with a well defined hazard rate.

6 Comparison to related literature

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 $\omega \in \{0, 1\}$, and

$$v(t, 0) = e^{-rt}, v(t, 1) = -ce^{-rt}.$$

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.⁵ Private signals affect only the relative likelihood of these two cases. This implies that no player ever wants to stop at any stage at some $t > 0$ conditional on no other investments within $(t - \varepsilon, t)$, since otherwise 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 real time zero, 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.

⁵We interpret here $\omega = 1$ to denote the case were it is never optimal to invest and $t = 1$ corresponds to the decision to never stop. To make the correspondence between the papers exact, postponing investment forever should actually give payoff zero in both states. This could be easily achieved by letting payoff be $v(t, 0) = e^{-r\frac{t}{1-t}}$, $v(t, 1) = -ce^{-r\frac{t}{1-t}}$, and interpreting the choice t as stopping at time $x = \frac{t}{1-t}$.

7 Conclusions

There are a number of directions where the analysis in this paper should be extended. First, considering our leading application, investment under uncertainty, one may view as quite extreme the modeling approach where nothing is learnt about the optimal investment time during the game from other sources than the behavior of the other players. Indeed, exogenous and gradually resolving uncertainty on the payoff of investment plays an important role in the literature on real options. The only reason why we have not included this element in our model is that we believe it would distract attention away from our main focus. Our paper can easily be extended to cover the case where the profitability of the investment depends on an exogenous (and stochastic) state variable p in addition to the private information about common market state ω . In this formulation, the stage game is a stochastic stopping game, where the players pick a Markovian strategy for optimal stopping. With our modeling assumptions this is equivalent to selecting a threshold value $p_i(\theta_i)$ for the exogenous state variable at which to stop, conditional on the signal. The stage ends at the first moment when the threshold value of some player is hit. With our monotonicity assumptions, there is an equilibrium with stopping thresholds that are monotonic in signals, hence a player with a low signal stops earlier than a player with a high signal. The substance of our results would carry over as such to this extended model.

The analytical simplicity of the model also makes it worthwhile to consider some other 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 discovered 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.

8 Appendix

Proof of Proposition 1. The monotonicity of $\tau^*(h^k, \theta)$ follows directly from MLRP and the (log-)supermodularity of v .

Denote by $\hat{\tau}(h^k; \theta)$ the optimal (unconstrained) stopping time based on the public history h^k and the knowledge that the lowest signal amongst the players remaining in the game after history h^k is θ :

$$\hat{\tau}(h^k; \theta) := \min \left(\arg \max_{t \geq 0} \mathbb{E} [v(t, \omega) \mid h^k, \theta_{\min}^k = \theta] \cdot \right) \quad (13)$$

The difference between $\hat{\tau}(h^k; \theta)$ and $\tau^*(h^k, \theta)$ defined in (3) is that the former measures real time from the beginning of the game and the latter measures time from the beginning of stage k . The relationship between the two is:

$$\tau^*(h^k, \theta) = \max(0, \hat{\tau}(h^k; \theta) - T^k). \quad (14)$$

Consider an arbitrary stage $k - 1$. The highest type that stops during that stage is $\underline{\theta}^k$, and therefore by (3) the duration of stage $k - 1$ is given by

$$\tau^*(h^{k-1}, \underline{\theta}^k) = T^k - T^{k-1}. \quad (15)$$

It follows from (14) that

$$\hat{\tau}(h^{k-1}; \theta) \leq T^{k-1} + \tau^*(h^{k-1}, \theta), \quad (16)$$

and combining (15) and (16), we have

$$\hat{\tau}(h^{k-1}; \underline{\theta}^k) \leq T^k.$$

Consider next stage k . We have $h^k = h^{k-1} \cup (t^k, \mathcal{S}^{k-1})$, where \mathcal{S}^{k-1} consists of players with signals in $(\underline{\theta}^{k-1}, \underline{\theta}^k)$. Therefore, it follows from MLRP and the (log-)supermodularity of v that

$$\hat{\tau}(h^k; \underline{\theta}^k) < \hat{\tau}(h^{k-1}; \underline{\theta}^k) \leq T^k.$$

By the continuity of signal densities, we then have

$$\widehat{\tau}(h^k; \underline{\theta}^k + \varepsilon) < T^k$$

for some $\varepsilon > 0$. But then from (14), we have

$$\tau^*(h^k; \underline{\theta}^k + \varepsilon) = 0,$$

and the result follows from the monotonicity of $\tau^*(h^k; \theta)$. ■

Proof of Theorem 1. The proof is done by using the one-shot deviation principle. We assume that all players $j \neq i$ use the strategy $\tau^*(h^k, \theta)$ after all histories h^k . We also assume that player i uses $\tau^*(h^{k'}, \theta)$ for all $k' > k$. We show that under this assumption, $\tau^*(h^k, \theta)$ is optimal for i after all histories h^k .

We show this in two steps. In the first step, we consider the optimal stopping problem of player i with signal θ after history stage h^k in an auxiliary problem where we require that $\tau(h^{k+1}, \theta) = 0$ after all h^{k+1} where i remains active. We show that in this auxiliary problem, $\tau^*(h^k, \theta)$ is an optimal stopping time. In the second step, we show that $\tau^*(h^k, \theta)$ is optimal also when i chooses $\tau^*(h^{k+1}, \theta)$ in stage $k + 1$.

We define the auxiliary problem as follows. The expected payoff of player i with signal θ when she chooses $t \geq 0$ in stage k is given by:

$$V^k(t, \theta) := \mathbb{E} [v(t^- + T^k, \omega) | h^k, \theta_i = \theta, t^-], \text{ where}$$

$$t^- = \min \left(t, \min_{j \neq i} \tau^*(h^k, \theta_j) \right).$$

The problem is to choose a t that maximizes $V^k(t, \theta)$. Since $v(t, \omega)$ is continuous in t , it follows that also $V^k(\cdot, \theta) : [0, \infty) \cup \infty \rightarrow \mathbb{R}$ is a continuous function.

Let \mathbb{T}^k the set of all t at which $\tau^*(h^k, \theta) = t$ for some θ .

$$\mathbb{T}^k := \left\{ t \geq 0 : \exists \theta \in [0, \bar{\theta}) \text{ s.t. } \lim_{\theta' \uparrow \theta} \tau^*(h^k, \theta') = \lim_{\theta' \downarrow \theta} \tau^*(h^k, \theta') = t \right\}.$$

Define also

$$\mathbb{T}_-^k := \mathbb{T}^k \cap [0, \tau^*(h^k, \theta)) \quad \text{and} \quad \mathbb{T}_+^k := \mathbb{T}^k \cap (\tau^*(h^k, \theta), \infty].$$

We show first that

$$\begin{aligned} \frac{\partial V^k(t, \theta)}{\partial t} &\geq 0 \text{ for } t \in \mathbb{T}_-^k \text{ and} \\ \frac{\partial V^k(t, \theta)}{\partial t} &\leq 0 \text{ for } t \in \mathbb{T}_+^k. \end{aligned}$$

The time derivative of $V^k(t, \theta)$ at $t \in \mathbb{T}^k$ can be written as:

$$\frac{\partial V^k(t, \theta)}{\partial t} = \Pr(\theta_{-i}^{\min} > \theta^{*k}(t)) \cdot \frac{\partial \mathbb{E}[v(t, \omega) \mid h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta]}{\partial t}, \quad (17)$$

where θ_{-i}^{\min} is the smallest signal amongst players other than i :

$$\theta_{-i}^{\min} := \min_{j \in \mathcal{N}^k \setminus i} \theta_j.$$

Note that when $t \in \mathbb{T}^k$, t is the optimal stopping time for type $\theta^{*k}(t)$.

Therefore, by optimality,

$$\begin{aligned} \frac{\partial V^k(t, \theta^{*k}(t))}{\partial t} &= \Pr(\theta_{-i}^{\min} > \theta^{*k}(t)) \cdot \frac{\partial \mathbb{E}[v(t, \omega) \mid h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta^{*k}(t)]}{\partial t} \\ &= \Pr(\theta_{-i}^{\min} > \theta^{*k}(t)) \cdot 0 = 0. \end{aligned} \quad (18)$$

Consider next an arbitrary posterior belief $p(\omega)$, and define $V_{p(\omega)}(t, \theta)$ as the expected value of stopping at t , when posterior is calculated based upon $p(\omega)$ and signal θ :

$$V_{p(\omega)}(t, \theta) := \mathbb{E}[v(t, \omega) \mid p(\omega), \theta_i = \theta].$$

Since we have assumed that $v(t, \omega)$ is differentiable in t for all $\omega \in \Omega$, $V_{p(\omega)}(t, \theta)$ must also be continuously differentiable. Moreover, it is straightforward to show that the supermodularity (or logsupermodularity) of $v(t, \omega)$ and MLRP together imply that

$$\frac{\partial V_{p(\omega)}(t, \theta)}{\partial t} \geq 0 \implies \frac{\partial V_{p(\omega)}(t, \theta')}{\partial t} \geq 0 \text{ for all } \theta' \in (\theta, \bar{\theta}), \quad (19)$$

$$\frac{\partial V_{p(\omega)}(t, \theta)}{\partial t} \leq 0 \implies \frac{\partial V_{p(\omega)}(t, \theta')}{\partial t} \leq 0 \text{ for all } \theta' \in [0, \theta]. \quad (20)$$

We apply these two equations to determine the sign of (17). If $t < (>) \tau^*(h^k, \theta)$, then we have $\theta > (<) \theta^{*k}(t)$. Combining (18) and (19) gives us

$$\frac{\partial V^k(t, \theta)}{\partial t} \geq (\leq) 0.$$

Next, consider the case where $t \notin \mathbb{T}^k$. Then, we must have

$$\tau^- \leq t \leq \tau^+,$$

where

$$\tau^- := \lim_{\theta \uparrow \theta^{*k}(t)} \tau^{*k}(\theta) < \tau^+ := \lim_{\theta \downarrow \theta^{*k}(t)} \tau^*(h^k, \theta).$$

If $t < \tau^{*k}(\theta)$, then $\theta > \theta^{*k}(t)$, and MLRP implies that

$$V^k(t, \theta) < V^k(\tau^+, \theta).$$

Similarly, if $t > \tau^*(h^k, \theta)$, then $\theta < \theta^{*k}(t)$, and MLRP implies that

$$V^k(t, \theta) < V^k(\tau^-, \theta).$$

We have now shown the following. First, $V^k(t, \theta)$ is increasing in t in the set \mathbb{T}_-^k and decreasing in \mathbb{T}_+^k . Second, when $t < \tau^*(h^k, \theta)$ and $t \notin \mathbb{T}_-^k$, we have

$$V^k(t, \theta) < V^k(\inf \{t' > t \mid t' \in \mathbb{T}_-^k\}, \theta),$$

and when $t > \tau^*(h^k, \theta)$ and $t \notin \mathbb{T}_+^k$, we have

$$V^k(t, \theta) < V^k(\sup \{t' < t \mid t' \in \mathbb{T}_+^k\}, \theta).$$

Since $V^k(t, \theta)$ is continuous in t , these results imply that:

$$\tau^*(h^k, \theta) \in \arg \max_t V^k(t, \theta).$$

We argue next that $\tau^*(h^k, \theta)$ remains optimal in stage k when i plays $\tau^*(h^{k+1}, \theta)$ in stage $k+1$. To see this, note that relaxing the constraint $\tau = 0$ in stage $k+1$ can only increase the optimal stopping time in stage k

(since it makes the continuation value in stage $k + 1$ larger). Therefore, it is immediate that our conclusion according to which $\tau^*(h^k, \theta)$ is preferred to all $t < \tau^*(h^k, \theta)$ continues to hold when we let i play $\tau^*(h^{k+1}, \theta)$ in stage $k + 1$.

On the other hand, we know from Proposition 1 that $\tau^*(h^{k+1}, \theta) = 0$ for all $\theta \leq \underline{\theta}^{k+1}$. In particular, this means that if stage k ends at time $t^k \geq \tau^*(h^k, \theta)$, we have $\theta \leq \underline{\theta}^{k+1}$ and i will in any case choose $\tau^*(h^{k+1}, \theta) = 0$. Therefore, for $t^k \geq \tau^*(h^k, \theta)$ the restriction $\tau(h^{k+1}, \theta_i) = 0$ is irrelevant because it is optimal to choose $\tau^*(h^{k+1}, \theta) = 0$.

We have now shown that if all players $j \neq i$ play $\tau^*(h^{k'}, \theta_{-i})$ in all stages $k' = 0, 1, \dots$, and if $\tau^*(\theta)$ is optimal for i in all stages $k' > k$, then $\tau^*(h^{k'}, \theta)$ is optimal for i in stage k . Since $\tau^*(h^k, \theta)$ is clearly also optimal for i in a stage where she is the only player left in the game, the proof is complete by backward induction.

Proof of Proposition 2. For $n = 1$, this result is implied by Theorem 5 of Gnedenko (1943). ■

To extend the result to $n > 1$, assume that $[Y_1^N, \dots, Y_k^N] \xrightarrow{\mathcal{D}} [Y_1^\infty, \dots, Y_k^\infty]$ for some $k \geq 1$. Consider Y_{k+1}^{N+1} . Since the signals are statistically independent, $\frac{N}{N+1} Y_{k+1}^{N+1}$ has the same distribution as $\{Y_k^N \mid Y_k^N \geq Y_{k+1}^{N+1}\}$. Since $Y_k^N \xrightarrow{\mathcal{D}} Y_k^\infty$ by assumption, we know from the properties of the exponential distribution that also Y_{k+1}^{N+1} converges to an independent exponentially distributed variable. Therefore, $Y_{k+1}^N \xrightarrow{\mathcal{D}} Y_{k+1}^\infty$. The result now follows from induction. ■

Proof of Lemma 1. Note first that $U(\cdot \mid z_n)$ is a continuous function on a compact set and therefore by Weierstrass' theorem $t_n(z_n)$ is non-empty. For the uniqueness, we use Theorem 1 in Araujo & Mas-Colell (1978). To this effect, we note that $v(t, \omega)$ is continuously differentiable in t for almost every (t, ω) and $\pi(\omega \mid z_n)$ is continuously differentiable in z_n . Therefore $U(t \mid z_n)$ is continuously differentiable in (t, z_n) . Furthermore, MLRP and the (log)supermodularity of $v(t, \omega)$ imply that for t and $t' \neq t$ such that $U(t \mid$

$z_n) = U(t' | z_n)$, we have:

$$\frac{\partial (U(t | z_n) - U(t' | z_n))}{\partial z_n} \neq 0.$$

Hence the conditions for Theorem 1 in Araujo & Mas-Colell (1978) are satisfied and the claim is proved. ■

Proof of Lemma 2. Consider the sequence $\{t_n^N(y_1, \dots, y_n)\}_{N=1}^\infty$. Since

$$\int_{\Omega} \pi^N(\omega | (y_1, \dots, y_n)) d\omega = 1 \text{ for all } N,$$

and $v(t, \omega)$ is bounded in ω , Corollary 2 implies that for every t ,

$$\lim_{N \rightarrow \infty} U^N(t | (y_1, \dots, y_n)) = U(t | y_n). \quad (21)$$

Moreover, since $v(t, \omega)$ is differentiable in t and this derivative is bounded for all ω , we have

$$\lim_{N \rightarrow \infty} \frac{\partial U^N(t | (y_1, \dots, y_n))}{\partial t} = \frac{\partial U(t | y_n)}{\partial t},$$

and therefore the convergence in equation (21) is uniform. Since $U^N(t | (y_1, \dots, y_n))$ is a continuous function on a compact set for all N , it has a maximum value. Uniform convergence then implies that

$$\lim_{N \rightarrow \infty} \left(\max_t (U^N(t | (y_1, \dots, y_n))) \right) = \max_t U(t | y_n).$$

Take any sequence $\{t_n^N\}_{N=1}^\infty$ such that $t_n^N \in \arg \max U^N(t | (y_1, \dots, y_n))$ for every N . Since $U^N(t | (y_1, \dots, y_n))$ converges uniformly to $U(t | y_n)$, and the latter has a unique maximizer $t_n(z_n)$ by Lemma 1, we have

$$t_n^N \rightarrow t_n(z_n).$$

The proof is identical for the sequence $\{t_n^N(z_n)\}_{N=1}^\infty$. ■

Proof of Proposition 3. Fix n and (y_1, \dots, y_n) . Call the player with the i^{th} lowest signal player i . Her normalized signal is y_i . Consider her information

at the time of stopping. By (3), she conditions on all the other remaining players having a signal higher than hers. Since the informative equilibrium is monotonic, all the players that have signals above her signal are active. Therefore, i conditions on her signal being the k^{th} lowest, where we must have $k \leq i$. It then follows from Lemma 2 that when $N \rightarrow \infty$, the optimal stopping time of i conditional on her information at the time of stopping converges to $t_k(z_i)$, where $k \leq i$. By MLRP and (log)supermodularity of v , we have $t_k(z_i) \geq t_i(z_i)$, and therefore,

$$\lim_{N \rightarrow \infty} T_i^N(y_1, \dots, y_i) \geq t_i(z_i). \quad (22)$$

Assume next that

$$\lim_{N \rightarrow \infty} T_i^N(y_1, \dots, y_i) > \lim_{N \rightarrow \infty} T_{i-1}^N(y_1, \dots, y_i). \quad (23)$$

This is the case, where player i stops at time $t^k > 0$ in some stage k (for N high enough). This means that i has the lowest signal among the active players at the time of stopping so that she correctly conditions on having the i^{th} lowest signal. Since her conditioning is correct, Lemma 2 implies that

$$\lim_{N \rightarrow \infty} T_i^N(y_1, \dots, y_i) = t_i(z_i). \quad (24)$$

Combining equations (22) - (24), we have

$$\lim_{N \rightarrow \infty} T_i^N(y_1, \dots, y_i) = \max \left[\lim_{N \rightarrow \infty} T_{i-1}^N(y_1, \dots, y_i), t_i(z_i) \right].$$

By the same arguments, we have

$$\lim_{N \rightarrow \infty} T_1^N(y_1) = t_1(z_1) = T_1(y_1).$$

Therefore, it follows by induction that for all $i = 1, \dots, n$

$$\lim_{N \rightarrow \infty} T_i^N(y_1, \dots, y_i) = \max [T_{i-1}(y_1, \dots, y_i), t_i(z_i)],$$

which is equivalent to

$$\lim_{N \rightarrow \infty} T_i^N(y_1, \dots, y_i) = \max_{n'=1, \dots, n} t_{n'}(z_{n'}).$$

■

Proof of Theorem 2. In a symmetric equilibrium, no information is transmitted before the first exit. By the monotonicity of the equilibrium strategies, a lower bound for all exit times and hence also for $T^N(\omega)$ for all N is $t(0)$. The fact that Y_1 has full support on the positive real line implies immediately that the distribution of the first exit time has support $[t(0), \bar{\omega}]$.

By Corollary 3, we know that $T_n^N \xrightarrow{D} T_n$. Therefore it is sufficient to consider the optimal decisions when players know their position in the ordering of the signals. Suppose a signal $\theta > 0$ is the $\alpha(N)^{\text{th}}$ lowest amongst N signals. Let $\alpha = \frac{\alpha(N)}{N}$. Then as $N \rightarrow \infty$ the weak law of large numbers implies that the true state ω can be found from the equation

$$G(\theta, \omega) = \alpha.$$

By our MLRP assumption, this is sufficient to identify ω . Therefore all players with signals $\theta > 0$ know the state when they invest (in the limit as $N \rightarrow \infty$). Therefore $\limPr\{T^N(\omega) \geq \omega\} = 1$. This establishes part 1. of the theorem.

The same logic yields immediately

$$\lim_{N \rightarrow \infty} \Pr\{|T^N(\omega, \theta) - T^N(\omega, \theta')| < \varepsilon\} = 1 \text{ for all } \varepsilon > 0,$$

as claimed in part 2 of the Theorem. ■

References

- Araujo, Aloisio & Andreu Mas-Colell. 1978. “Notes on the smoothing of aggregate demand.” *Journal of Mathematical Economics* 5(2):113 – 127.
- Banerjee, A.V. 1992. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics* 107:797–817.

- Bikhchandani, S., D. Hirshleifer & I. Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100:992–1026.
- 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.
- Gnedenko, B. 1943. "Sur La Distribution Limite Du Terme Maximum D'Une Serie Aleatoire." *The Annals of Mathematics* 44(3):pp. 423–453.
- Grenadier, Steven. 1999. "Information revelation through option exercise." *Review of Financial Studies* 12(1):95–129.
URL: <http://rfs.oxfordjournals.org/content/12/1/95.abstract>
- Kremer, I. 2002. "Information Aggregation in Common Value Auctions." *Econometrica* 70:1675–1682.
- Levin, Dan & James Peck. 2008. "Investment dynamics with common and private values." *Journal of Economic Theory* 143(1):114 – 139.
- Moscarini, Giuseppe & Francesco Squintani. 2010. "Competitive experimentation with private information: The survivor's curse." *Journal of Economic Theory* 145(2):639 – 660.
- Murto, Pauli & Juuso Välimäki. 2011. "Learning and information aggregation in an exit game." *Review of Economic Studies* forthcoming.
- Pesendorfer, W. & J. Swinkels. 1997. "The Loser's Curse and Information Aggregation in Common Value Auctions." *Econometrica* 65(6):1247–1282.




Figure 1: Payoffs as functions of signal precision in the quadratic binary example.




Figure 2: Probability distribution of the stopping time of the last player with various signal precisions.