Delay and Information Aggregation in Stopping Games with Private Information

Pauli Murto[*] Juuso Välimäki [†]

August 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 is one of observational learning where communication between the players is not considered.

For concreteness, one may interpret the stopping decision as an irreversible investment decision as in the literature on real options. Since the payoff relevant parameter is common to all players, the equilibrium investments are complementary. Delayed investment by other firms indicates less favorable conditions for early investment whereas early investment by other firms encourages the others to follow. To put it simply, the first firm to invest must always worry about the fact that others have not yet invested. This creates a hurdle to investment. But once this hurdle is cleared and some firms invest, it is likely

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

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

that the other firms realize that they are already too late. This starts an investment wave.

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. iv) Almost all players benefit from observational learning.

In our model, the first-best time to invest is common to all players and depends on a single state variable ω. Since all the players have information on ω, the observed actions reflect 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 investment 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 decision for each player in the most informative symmetric equilibrium of the game. In this equilibrium, a player always calculates her payoffs 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 invest immediately regardless of their signals.

We allow the players to react quickly to each other's decisions. In order to avoid complicated limiting procedures, we model the dynamic game as a multi-stage game with continuous action sets. 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. Quick reactions to stopping decisions are captured by allowing immediate stopping in the next stage. It is well-know that for some stopping games with payoff externalities, the existence of a stage game equilibrium is problematic in the continuous action variable case. In our game, all externalities are informational and no such difficulties arise.[1]

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.

If N is large, the weak law of large numbers guarantees that the number of players with a private signal below any arbitrary value θ identifies the state ω (approximately) accurately. 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. This asymmetry that arises in most timing games explains the delays in our model.

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

[1] An early example of such existence problems appears in Fudenberg & Tirole (1985). With private information, equilibrium existence is less problematic than in complete information settings. This can be easily demonstrated in two-player games with a first mover advantage.

information is revealed inefficiently late.

The most closely related paper is the investment model by Chamley & Gale (1994) where in contrast to our paper, it is optimal to invest either immediately or never.[2] We allow for a more general payoff structure where the state of nature determines 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.[3] 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 is revealed 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 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 presents some extensions of the basic model and compares our results to the most closely related literature. Section 7 concludes.

[2] See 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.

[3] Section 6 discusses in more detail the relationship between these papers.

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 subset of the extended positive real line $\Omega \subseteq [0, \infty]$, and can be either finite or infinite.[4] The players share a common prior $p^0(\omega)$ on Ω. The players choose 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 the following:

Assumption 1 *The payoff function $v(t, \omega)$ is twice differentiable in t almost everywhere, and for each ω, there is a unique t that maximizes $v(t, \omega)$. Furthermore, $v(t, \omega)$ is either strictly supermodular, or strictly log-supermodular in (t, ω).*

The key implication of the assumption of strict (log-)supermodularity is that the unique maximizer of $v(t, \omega)$ must be strictly increasing in ω. Examples include: i) Quadratic loss relative to optimal time $\omega : v(t, \omega) = -(t - \omega)^2$. ii) Discounted model of costly investment where the market becomes profitable at random time $\omega : v(t, \omega) = e^{-r \max\{t,\omega\}} - Ce^{-rt}$, where $0 < C < 1$ is a parameter. iii) "Now or never": a special case of ii) with state space $\Omega = \{0, \infty\}$. iv) Discounted costly investment in a market growing at rate $\alpha < r$: $v(t, \omega) = e^{-rt}(e^{\alpha t} - \omega)$.[5]

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 $G(\theta \mid \omega)$ and corresponding densities $g(\theta \mid \omega)$ are well defined and have full support for all ω. We also assume that for all ω, $G(\theta \mid \omega)$ is continuous (i.e., there are no mass points) and $g(\theta \mid \omega)$ has at most a finite number of points of discontinuity and is continuous at $\theta = 0$.

The signals in the support of the signal distribution satisfy monotone likelihood ratio property (MLRP):

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

$$\frac{g(\theta' \mid \omega')}{g(\theta \mid \omega')} \geq \frac{g(\theta' \mid \omega)}{g(\theta \mid \omega)}. \tag{1}$$

[4] We also assume that $v(t, \omega)$ satisfies continuity at infinity in both t and ω to ensure the existence of optimal decisions.

[5] A variant of this model with a stochastic state variable will be discussed in Section 6.3.

Assumptions 1 and 2 together allow 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 $0 < g(0|\omega) < \infty$ 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 $\tau(h^0, \theta) = (\tau_1(h^0, \theta_1), ..., \tau_N(h^0, \theta_N))$ 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:

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

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. Stage k starts at the point in time t^{k-1} where the previous stage ended. The players that have not yet invested choose an investment time $\tau_i(h^k, \theta_i) \geq t^{k-1}$. 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

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

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 $Q^k := \sum_{i=0}^{k} S^k$.

A pure behavior strategy for stage k is a function

$$\tau_i^k : H^k \times \Theta \to [t^{k-1}, \infty],$$

and we also define the strategy $\tau_i(h, \theta)$ on the set of all histories by:

$$\tau_i(h, \theta) = \tau_i^k(h, \theta) \text{ whenever } h \in H^k.$$

The players maximize their expected payoff. A strategy profile $\tau = (\tau_i, ..., \tau_N)$ 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 $\tau_{-i.}$

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, $\tau_i\left(h^k, \theta\right)$ 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 $\left(\underline{\theta}^k, \bar{\theta}\right)$, where:

$$\underline{\theta}^k := \sup\left\{\theta \left| \tau(h^{k-1}, \theta) = t^{k-1}\right.\right\}. \tag{2}$$

3.1 Informative Equilibrium

We now characterize the symmetric equilibrium that maximizes information transmission in the set of symmetric monotone pure strategy equilibria. Theorem 1 below 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:

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

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:*

$$\tau^*\left(h^k, \theta\right) := \min\left(\arg\max_{t \geq t^{k-1}} \mathbb{E}\left[v\left(t, \omega\right) \left| h^k, \theta^k_{\min} = \theta\right.\right]\right). \tag{3}$$

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

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

Proof. Proposition 1 is proved in the Appendix. ■

The equilibrium stopping strategy $\tau^*\left(h^k, \theta\right)$ defines a time-dependent cutoff signal $\theta^{*k}(t)$ for all $t \geq t^{k-1}$:

$$\theta^{*k}(t) := \sup\left\{\theta \left| \tau^*\left(h^k, \theta\right) \leq t\right.\right\}. \tag{4}$$

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, $\theta^{*k}\left(t^{k-1}\right) > \underline{\theta}^k$ for all stages except possibly the first one. This means that all the players with a signal in the interval $\left(\underline{\theta}^k, \theta^{*k}\left(t^{k-1}\right)\right)$ 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 the cutoff signal $\theta^{*k}(t)$ (i.e. the lower bound of the signals of the existing players) moves upwards as time goes by. 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 stage k ends at $t^k > t^{k-1}$, the remaining players learn that one of the other active players in stage k has a signal at the lower bound $\theta^{*k}\left(t^k\right)$. By MLRP and the (log)supermodularity of v, the expected value from staying in the game falls by a discrete amount. This means that the cutoff type moves discretely upwards and explains why $\theta^{*k+1}\left(t^k\right) > \theta^{*k}\left(t^k\right) = \underline{\theta}^{k+1}$. As a result, each new stage begins with a positive probability of immediate further exits. If at least one player stops so that $t^{k+1} = t^k$, 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 $t^{k'} = t^{k'-1}$ 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 $t^k = t^{k-1}$. 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 $t = t^{k-1}$ solves that problem for all signal types remaining in the game, i.e., if

$$t^{k-1} \in \arg\max_{t \geq t^{k-1}} \mathbb{E}\left[v\left(t, \omega\right) \middle| h^k, \theta_i = \theta \right] \text{ for all } \theta \geq \underline{\theta}^k,$$

then an uninformative equilibrium may exist. If all players stop at $t = t^{k-1}$ then they learn nothing from each other. If they learn nothing from each other, then $t = t^{k-1}$ is their optimal action.

It should be noted that some equilibria where all the players stop immediately satisfy our criteria for informative equilibrium. If $\tau^*(h^k, \theta) = t^{k-1}$ 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 $t = \infty$ (which, in such a case, is to be interpreted as delaying infinitely). See discussion of such a case in Section 6.1.

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 $\tau^*\left(h^k,\theta\right)$ 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 analyze the limiting properties of the model as we increase the number of players towards infinity. Since the informative equilibrium strategy is monotonic in signals, the players stop in the ascending order of their signals. Therefore, given the game with N players, the time instant at which the n:th player stops the game is a function of the n lowest signal realizations amongst the players, and we can write it as

$$T_n^N\left(\widetilde{\theta}_1^N,...,\widetilde{\theta}_n^N\right),$$

where $\widetilde{\theta}_i^N$ denotes the i^{th} order statistic in the game with N players:

$$\widetilde{\theta}_i^N := \min\left\{\theta\in\left[0,\overline{\theta}\right) \mid \#\left\{j\in\mathcal{N}\mid\theta_j\leq\theta\right\}=i\right\}. \tag{5}$$

We start with a simple statistical observation regarding the distribution of order statistics in large samples (Section 4.1). Using this result, we then derive the limiting distribution for the n first stopping moments in the informative equilibrium (Section 4.2). Finally, we present a theorem that characterizes the equilibrium stopping times of (almost) all the players in the large-game limit (Section 4.3).

4.1 Information in equilibrium

It is clear that if we increase N towards infinity while keeping n fixed, the n lowest order statistics $\widetilde{\theta}_1^N,...,\widetilde{\theta}_n^N$ converge to the lower bound 0 of the signal distribution in probability. Therefore, we scale the order statistics by the number of players:

$$Z_i^N := \widetilde{\theta}_i^N\cdot N. \tag{6}$$

Since Z_i^N is a deterministic function of $\widetilde{\theta}_i^N$, it has the same information content as $\widetilde{\theta}_i^N$. In the next proposition we record a well known statistical result according to which Z_i^N

converge to non-degenerate random variables. This limit distribution, therefore, captures the information content of $\widetilde{\theta}_n^N$ in the limit.

Proposition 2 *For all $n \in \mathbb{N}$, the vector $\left[Z_1^N, Z_2^N - Z_1^N, ..., Z_n^N - Z_{n-1}^N\right]$ converges in distribution to a vector of n independent exponentially distributed random variables with parameter $g(0 \mid \omega)$. That is,*

$$\begin{aligned} & \lim_{N\to\infty} \Pr\left(Z_1^N \leq x_1, Z_2^N - Z_1^N \leq x_2, ..., Z_n^N - Z_{n-1}^N \leq x_n\right) \\ = \; & e^{-g(0|\omega)\cdot x_1} \cdot ... \cdot e^{-g(0|\omega)\cdot x_n}. \end{aligned}$$

Proof. In the Appendix. ■

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

Note that $Z_n^N = Z_1^N + \left(Z_2^N - Z_1^N\right) + ... + \left(Z_n^N - Z_{n-1}^N\right)$, and therefore Z_n^N converges to a sum of independent exponentially distributed random variables, which means that its limiting distribution is Gamma:

Corollary 1 *For all n,*

$$Z_n^N \xrightarrow{\mathcal{D}} \Gamma(n, g(0 \mid \omega)),$$

where $\Gamma(n, g(0 \mid \omega))$ denotes gamma distribution with parameters n and $g(0 \mid \omega)$.

We have now seen that when $N \to \infty$, observing the n lowest order statistics is equivalent to observing n independent exponentially distributed random variables. Since exponential distributions are 'memoryless', this means that observing only the n^{th} order statistic $\widetilde{\theta}_n^N$ is informationally equivalent to observing all order statistics up to n. To see this important fact formally, denote by $\pi(\omega \mid (z_1, ..., z_n))$ the posterior probability of an arbitrary element $\omega \in \Omega$ based on a realization $(z_1, z_2 - z_1, ..., z_n - z_{n-1})$ of independent exponential variables, and let $\pi(\omega \mid z_n)$ denote the corresponding posterior probability based on the sample that contains only z_n, the sum of the previous sample. Bayes' rule and simple algebra show that these posteriors are equal:

$$\begin{aligned} \pi(\omega \mid (z_1, ..., z_n)) &= \frac{\pi^0(\omega) \cdot \prod_{i=1}^{n} g(0 \mid \omega) e^{-g(0|\omega)(z_i - z_{i-1})}}{\int_\Omega \pi^0(\omega') \cdot \prod_{i=1}^{n} g(0 \mid \omega') g(0 \mid \omega') e^{-g(0|\omega')(z_i - z_{i-1})} d\omega'} \qquad (7) \\ &= \frac{\pi^0(\omega) \cdot (g(0 \mid \omega))^n e^{-g(0|\omega) z_n}}{\int_\Omega \pi^0(\omega') \cdot (g(0 \mid \omega'))^n e^{-g(0|\omega') z_n} d\omega'} = \pi(\omega \mid z_n). \end{aligned}$$

In the finite model (away from the limit $N \to \infty$), the posterior $\pi^N(\omega \mid (z_1, ..., z_n))$ based on a sample $Z_1^N = z_1, ..., Z_n^N = z_n$ generally differs from the posterior $\pi^N(\omega \mid z_n)$ that is based only on $Z_n^N = z_n$. Nevertheless, Bayes' rule is continuous in the limit as $N \to \infty$ in $(z_1, ..., z_n)$ since we assume $g(\cdot \mid \omega)$ to be continuous at $\theta = 0$ for all ω. Therefore, Proposition 2 implies that both $\pi^N(\omega \mid (z_1, ..., z_n))$ and $\pi^N(\omega \mid z_n)$ converge to the posterior $\pi(\omega \mid z_n)$ for all ω and $(z_1, ..., z_n)$ as $N \to \infty$. We summarize this discussion in the following Corollary.

Corollary 2 *For a fixed sample of normalized order statistics* $(z_1, ..., z_n)$*,*

$$\lim_{N\to\infty} \pi^N(\omega \mid (z_1, ..., z_n)) = \lim_{N\to\infty} \pi^N(\omega \mid z_n) = \pi(\omega \mid z_n) \text{ for all } \omega.$$

More generally, a player may have some, but not perfect, information on $(z_1, ..., z_{n-1})$. Suppose that a player knows z_n, and in addition knows that each z_i, $i < n$, lies within some arbitrary interval A_i of the real line. Corollary 2 also means that

$$\lim_{N\to\infty} \pi^N(\omega \mid z_1 \in A_1, ..., z_{n-1} \in A_{n-1}, z_n) = \pi(\omega \mid z_n).$$

This observation plays a key role in our analysis. Suppose that player i has signal θ and that she has some information on the signals of those players that have stopped before her. In particular, by the monotonicity of the informative equilibrium strategy profile, she knows at the very least that those signals are all below θ. By Theorem 3, she would now choose the optimal stopping time conditional on her information on those 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 $N \to \infty$. Hence even if all signals were observable, the relevant conditioning event is still $Z_n^N = N\theta$ when $N \to \infty$. We now turn to the formalization of this reasoning.

4.2 Timing in Large Games

In this section, we link the equilibrium stopping decisions to the information contained in the order statistics. We show that when $N \to \infty$, the equilibrium path of the game can be approximated by a simple algorithm that samples sequentially the order statistics.

As a terminological matter, we use the term *unconstrained* stopping time to refer to the optimal element from the original action space $T = [0, \infty]$. Since in equilibrium the players cannot go backwards in time, we are ultimately interested in stopping times chosen from $\left[t^{k-1}, \infty\right]$ where t^{k-1} is the time at which the previous stage ended. We use the term *constrained* stopping time to refer to an optimal stopping time that is constrained to be weakly higher than some previously chosen stopping time. We use the term *limit model* to refer to the statistical properties of the order statistics in the limit $N \to \infty$, derived in the previous subsection.

We consider first the unconstrained stopping time in a hypothetical case, where a decision maker observes the n^{th} order statistic of the limit model. In the following Lemma we establish the uniqueness of the optimal solution to this problem for almost every realization z_n of Z_n.

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

$$t_n(z_n) := \arg\max_{t \in [0,\infty]} \int_\Omega v(t, \omega) \pi(\omega \mid z_n) d\omega. \tag{8}$$

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

$$\left(Z_1^N = z_1, ..., Z_n^N = z_n\right),$$

and let $t_n^N(z_1, ..., z_n)$ and $t_n^N(z_n)$ denote the unconstrained optimal stopping times, based on the whole sample $(z_1, ..., z_n)$ and sample z_n, respectively:

$$\begin{aligned} t_n^N(z_1, ..., z_n) &: = \arg\max_{t \in [0,\infty]} \int_\Omega v(t, \omega) \pi^N(\omega \mid (z_1, ..., z_n)) d\omega, \\ t_n^N(z_n) &: = \arg\max_{t \in [0,\infty]} \int_\Omega v(t, \omega) \pi^N(\omega \mid z_n) d\omega. \end{aligned}$$

Note that $t_n^N(y_1, ..., y_n)$ and $t_n^N(z_n)$ could in principle be sets. The next Lemma, which is based on Corollary 2 in the previous subsection, shows that they converge to $t_n(z_n)$, which is singleton for almost every z_n by Lemma 1.

Lemma 2 *For almost every $(z_1, ..., z_n)$,*

$$\lim_{N \to \infty} t_n^N(z_1, ..., z_n) = \lim_{N \to \infty} t_n^N(z_n) = t_n(z_n).$$

Proof. In the Appendix. ■

With this Lemma, we can relate the equilibrium stopping times to the stopping times of the limit model. Notice, however, that so far we have been discussing the unconstrained stopping times $t_n(z_n)$, which could be decreasing in n. Since the players cannot go backwards in time, the relevant constrained stopping time in the limit model for the player with n^{th} lowest signal is the maximum of $t_{n'}(z_{n'})$, $n' = 1, ..., n$:

$$\bar{t}_n(z_1, ..., z_n) := \max_{n'=1,...,n} t_{n'}(z_{n'}). \tag{9}$$

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 (9). Recall that $T_n^N\left(\widetilde{\theta}_1^N, ..., \widetilde{\theta}_n^N\right)$ denotes the real time at which the player with the n^{th} lowest signal stops in the informative equilibrium. We have:

Proposition 3 *For all n, and for almost every $(z_1, ..., z_n)$,*

$$\lim_{N\to\infty} T_n^N\left(\frac{z_1}{N}, ..., \frac{z_n}{N}\right) = \bar{t}_n(z_1, ..., z_n).$$

Proof. In the Appendix. ■

We end this section by relating the joint distribution of equilibrium stopping times to the stopping times of the limit model. Omitting the arguments, let $\left[T_1^N, ..., T_n^N\right]$ denote the vector that contains the random stopping times of the n first players to stop in the symmetric equilibrium. Corollary 3 below provides a simple algorithm for simulating equilibrium stopping times in the large-game limit: fix an arbitrary n, draw n realizations $(z_1, ..., z_n)$ from exponential distribution with parameter $g(0 \mid \omega)$, and compute $\bar{t}_1(z_1), ..., \bar{t}_n(z_1, ..., z_n)$ using (8) and (9).

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

$$\left[T_1^N, ..., T_n^N\right] \xrightarrow{\mathcal{D}} \left[\bar{t}_1(Z_1), ..., \bar{t}_n(Z_1, ..., Z_n)\right]$$

where $\bar{t}_i$ is a function defined by (8) and (9), and $Z_1, ..., Z_n$ are independent, exponentially distributed random variables with parameter $g(0 \mid \omega)$.

Proof. Direct consequence of Propositions 2 and 3. ■

4.3 Delay in Equilibrium

In this section, we characterize the real time behavior of (almost) all the players in the informative equilibrium when $N \to \infty$. Let $T^N(\theta, \omega)$ denote the random stopping 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

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

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

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

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

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

$$F(t \mid \omega) = \Pr\{T(\omega) \leq t\}.$$

The following Theorem characterizes the asymptotic behavior of the informative equilibrium as the number of players becomes large. We denote by $t\,(0)$ the optimal investment time of a player that decides based on signal $\theta = 0$ only, and we denote by $t^*\,(\omega)$ the first-best investment time for state ω:

$$t^*\,(\omega) := \arg \max_{t \in [0,\infty]} v\,(t, \omega)\,.$$

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

1. *For all $\theta > 0$,*

$$\lim_{N \to \infty} \Pr\{\left|T^N(\omega, \theta) - T^N(\omega)\right| < \varepsilon\} = 1 \textit{ for all } \varepsilon > 0.$$

2. *$F(t \mid \omega) = 0$ for all $t < \max\{t(0), t^*\,(\omega)\}$.*

3. *$F(t \mid \omega) < 1$ for all $t < \max \Omega$.*

Proof. In the Appendix. ■

Theorem 2 confirms the main properties of our model. Almost all the players stop (almost) simultaneously (Part 1 of the theorem), and this stopping moment is inefficiently late and random (Parts 2 and 3 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 limit model 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:

$$\frac{g(\theta \mid 0)}{g(\theta \mid 1)} = c_l \qquad \text{for all } 0 \leq \theta \leq \theta^*, \tag{10}$$

$$\frac{g(\theta \mid 0)}{g(\theta \mid 1)} = c_h \qquad \text{for all } \theta^* < \theta < \overline{\theta}, \tag{11}$$

where $c_l > c_h > 0$ and $\theta^* > 0$ are parameters. 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 the probability of getting a low (high) signal if $\omega = 0$ ($\omega = 1$) is given by a parameter $\alpha > 1/2$:

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

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

$$v(t, \omega) = -(t-\omega)^2. \tag{12}$$

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 from other players. First, suppose a player must choose the stopping time without a signal. Then she stops at $t = 1/2$ and 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 $\theta \le \theta^*$, 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 non-optimal 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 $V^I \downarrow V^0$.

Consider next the case with a large N. If 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 1, 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 1). 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 $t^*(\omega) = 1$, and Part 2 of Theorem 2 states that in the large game limit the last player stops at time $t = 1$. Part 1 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 $\tilde{T} \to 1$ in probability. Therefore, denoting the expected payoff conditional on state ω by V_ω, we have:

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

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

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

$$\tilde{V}_{\theta=l} = V^I = -\alpha(1-\alpha). \tag{13}$$

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

$$\widetilde{V}_{\theta=l} = \alpha \widetilde{V}_0 + (1-\alpha)\,\widetilde{V}_1. \tag{14}$$

Combining (13) and (14), and solving for $\widetilde{V}_0$ gives:

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

Therefore,

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

The final step is to observe that as $N \to \infty$, we have $\tilde{V} \to V^*$ in probability. This follows from Part 1 of Theorem 2: since the real stopping times of all signal types (expect zero-probability case $\theta = 0$) converge to the same instant, the deviation that we have considered will not affect the realized payoff in the large game limit. Therefore, as $N \to \infty$,

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

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, denoting by V_ω^* the equilibrium payoff conditional on state, it should be noted that $V_1^* \to 0$ and $V_0^* \to -(1-\alpha)$. 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 $\tilde{T} \to 1$ in probability if $\omega = 1$, and therefore, if it turns out that $\tilde{T} < 1$, then we 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 $N \to \infty$ 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 $\Pr\{\tilde{T} < 1 - \epsilon \,|\omega = 1\} \to 0$ as $N \to \infty$, we can write the conditional probabilities of the event $\{\tilde{T} \in [t, t + dt) \mid \tilde{T} \geq t\}$ as

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

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 >

It should be noted that the binary state-space makes this example quite special. With more than two states, we are not able to compute analytically the equilibrium payoffs or the probability distribution for the players' stopping times. Nevertheless, as explained in Section 4.2, it is easy to simulate the large-game limit for any model specification. As an illustration, we extend the example to ten states: $\omega \in \{0, \frac{1}{9}, \frac{2}{9}, ..., 1\}$ (the payoff is given by (12) as before so that $t^*(\omega) = \omega$). Since we simulate the model directly in the large-game limit, we only need to specify the signal distributions at the low end of the signal space, and we let

$$g(0 \mid \omega) = 1 - \alpha \left(\omega - \frac{1}{2}\right),$$

where $\alpha \in [0, 2)$ is a parameter that measures the precision of the signals. We use Monte-Carlo simulation to derive V_ω^* and $F(t \mid \omega)$ for all state values with two signal precisions: $\alpha = 1$ (precise signals) and $\alpha = 0.1$ (imprecise signals). Figure 3 shows V_ω^*. We see that V_ω^* is increasing in ω so that observational learning is especially beneficial in those states where first-best investment is late. Also, we see that V_ω^* is higher for $\alpha = 1$ so that the players benefit from more accurate signals.

< Figure 3 here >

Figure 4 shows $F(t \mid \omega)$ for all state values (upper panel with $\alpha = 1$, lower panel with $\alpha = 0.1$). This figure confirms the properties derived in Theorem 2: for any state realization, the players stop at a random time that is always later than the first-best time. Note that there is more delay with imprecise signals, which explains the higher payoffs with precise signals.

< Figure 4 here >

To summarize, this quadratic 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 $N \to \infty$. iv) The instant at which almost all the players invest arrives with a well defined hazard rate.

6 Discussion

6.1 Relation to Chamley and Gale (1994)

Our general results are quite different from the 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

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

This is the special case, where the optimal investment takes place either immediately or never. The private signals affect only the relative likelihood of these two cases. To see this formally, note that for any information that a player might have, the strategy defined in Theorem 1 is always a corner solution: either $\tau^*(h^t, \theta) = t^{k-1}$ or $\tau^*(h^t, \theta) = \infty$. In other words, as explained in Chamley & Gale (1994), no player ever stops in any stage at some $t > t^{k-1}$ conditional on no other investments within $(t - \varepsilon, t)$ since otherwise it would have been optimal to invest already at $t - \varepsilon$. As a result, a given stage k ends either immediately if at least one player stops at time $t^k = t^{k-1}$ or the stage continues forever. Since this holds for all stages, all investment in the game 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 the limit equilibrium in their model as the period length is reduced corresponds exactly to the informative equilibrium of this special case of our model.

6.2 Uninformed Investors

Suppose that there are N informed players and a random number of uninformed investors. For simplicity, one could assume that the uninformed investors arrive according to an exogenously given Poisson rate λ per unit of real time. Assuming that the players are anonymous, the statistical inference is changed only minimally relative to our current model. If $t^k > t^{k-1}$, then there is a positive probability that the stopping player is indeed uninformed. As a result, the remaining players update their beliefs less than in the main model.

In any stage where $t^k = t^{k-1}$, the player that stops is informed with probability 1. This conclusion follows from the fact that stage k has a real-time duration 0 and uninformed investors arrive at a bounded rate λ. Hence inference in such stages is identical to the main model and all the qualitative conclusions remain valid. It can be shown that for large games, the hazard rate with which the game ends is unchanged by the introduction of uninformed players as long as λ is bounded.

6.3 More General State Variables

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.

Our paper can easily be extended to cover the case where the profitability of the investment depends on an exogenous (and stochastic) state variable in addition to the private information about common market state ω. An example of such a formulation is:

$$\begin{aligned} v(t, \omega\,; x) &= e^{-rt}\left(x_t - \omega\right), \\ \frac{dx_t}{x_t} &= \alpha dt + \sigma dZ_t, \end{aligned}$$

where Z_t is a Brownian motion. Such investment problems have been studied extensively in the literature (see Dixit & Pindyck (1994) for a survey), and it is well known that the optimal investment time is the smallest t where x_t exceeds a threshold value $x(\omega)$. Hence the problem is reduced to a model with a single state x_t, and the optimal investment threshold for a known ω is strictly increasing in ω. The analysis of our paper would extend in a straightforward manner to this case: the informative equilibrium strategy would command a player with signal θ to choose an investment threshold $x^*\left(h^k, \theta\right)$ that is optimal conditional on θ being the lowest signal among the remaining players. By our assumption of MLRP of the signals, the equilibrium thresholds would always be increasing in θ. All of our results would have a natural analogue in this extended model, with the stochastic state variable x_t playing the role that the calender time t plays in the current paper.

7 Conclusions

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 $\tau^*\left(h^k,\theta\right)$ follows directly from MLRP and the (log-)supermodularity of v.

Denote by $\widehat{\tau}\left(h^k,\theta\right)$ 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:

$$\widehat{\tau}\left(h^k,\theta\right) := \min\left(\arg\max_{t\geq 0}\mathbb{E}\left[v\left(t,\omega\right)\left|h^k,\theta^k_{\min}=\theta\right.\right]\right). \tag{15}$$

The relationship between $\widehat{\tau}\left(h^k;\theta\right)$ and $\tau^*\left(h^k,\theta\right)$ defined in (3) is:

$$\tau^*\left(h^k,\theta\right) = \max\left(t^{k-1},\widehat{\tau}\left(h^k,\theta\right)\right). \tag{16}$$

Consider an arbitrary stage $k-1$. The highest type that stops during that stage is $\underline{\theta}^k$, and therefore by (16)

$$\widehat{\tau}\left(h^{k-1},\underline{\theta}^k\right) \leq \tau^*\left(h^{k-1},\underline{\theta}^k\right) = t^{k-1}. \tag{17}$$

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

$$\widehat{\tau}\left(h^k,\underline{\theta}^k\right) < \widehat{\tau}\left(h^{k-1},\underline{\theta}^k\right) \leq t^{k-1},$$

where the latter inequality follows from (17). By the continuity of signal densities, we then have

$$\widehat{\tau}\left(h^k,\underline{\theta}^k+\varepsilon\right) < t^{k-1}$$

for some $\varepsilon>0$. But then from (16), we have

$$\tau^*\left(h^k,\underline{\theta}^k+\varepsilon\right) = t^{k-1},$$

and the result follows from the monotonicity of $\tau^*\left(h^k,\theta\right)$ in θ. ■

Proof of Theorem 1. The proof uses the one-shot deviation principle. We assume that all players $j\neq i$ play according to $\tau^*\left(h,\theta\right)$ after all histories h. We consider an arbitrary history h^k and assume that player i uses $\tau^*\left(h^{k'},\theta\right)$ for all $k'>k$. We show that under this assumption, $\tau^*\left(h^k,\theta\right)$ is optimal for i.

We divide the proof in two steps. In the first step, we consider the optimal stopping problem of player i in an auxiliary problem where we impose $\tau_i(h^{k+1},\theta)=t^k$ for all h^{k+1}

where i remains active. That is, we force i to stop immediately if some other player $j \neq i$ stops first in stage k. In the second step, we show that $\tau^*\left(h^k, \theta\right)$ remains optimal when i chooses $\tau^*\left(h^{k+1}, \theta\right)$ in stage $k+1$.

Step 1: We define the auxiliary problem as follows. Consider player i with signal θ at an arbitrary stage k. Her expected payoff when she chooses $t \geq t^{k-1}$ in stage k is given by:

$$\begin{aligned} V^k(t, \theta) \quad &: \quad = \mathbb{E}\left[v\left(t^-, \omega\right) \middle| h^k, \theta_i = \theta, t^-\right], \text{ where} \\ t^- \quad &= \quad \min\left(t, \min_{j \neq i} \tau^*\left(h^k, \theta_j\right)\right). \end{aligned}$$

The problem is to choose a $t \geq t^{k-1}$ that maximizes $V^k(t, \theta)$. Since $v(t, \omega)$ is continuous in t, it follows that also $V^k(\cdot, \theta) : \left[t^{k-1}, \infty\right] \to \mathbb{R}$ is a continuous function. Our goal is to show that $V^k(t, \theta)$ defined above is maximized at $t = \tau^*\left(h^k, \theta\right)$.

Let $\mathbb{T}^k$ define the subset of $\left(t^{k-1}, \infty\right)$ that consists of those (interior) time instants that are chosen in equilibrium for some signal θ':

$$\mathbb{T}^k := \left\{t \in \left(t^{k-1}, \infty\right) : \exists \theta' \in \left[0, \bar{\theta}\right) \text{ s.t. } \lim_{\theta'' \uparrow \theta'} \tau^*\left(h^k, \theta''\right) = \lim_{\theta'' \downarrow \theta'} \tau^*\left(h^k, \theta''\right) = t\right\}.$$

The plan for accomplishing step 1 of the proof is as follows. First, we show that

$$\begin{aligned} \frac{\partial V^k(t, \theta)}{\partial t} \quad &\geq \quad 0 \text{ for } t \in \mathbb{T}^k \cap \left[0, \tau^*\left(h^k, \theta\right)\right) \text{ and} \\ \frac{\partial V^k(t, \theta)}{\partial t} \quad &\leq \quad 0 \text{ for } t \in \mathbb{T}^k \cap \left(\tau^*\left(h^k, \theta\right), \infty\right]. \end{aligned} \tag{18}$$

Second, we show that when $t^{k-1} < t < \tau^*\left(h^k, \theta\right)$ and $t \notin \mathbb{T}^k$, we have

$$V^k(t, \theta) \leq V^k\left(\inf\left\{t' > t \middle| t' \in \mathbb{T}^k\right\}, \theta\right), \tag{19}$$

and when $\tau^*\left(h^k, \theta\right) < t < \infty$ and $\tau \notin \mathbb{T}^k$, we have

$$V^k(t, \theta) \leq V^k\left(\sup\left\{t' < t \middle| t' \in \mathbb{T}^k\right\}, \theta\right). \tag{20}$$

Since $V^k(t, \theta)$ is continuous in t, (18) - (20) imply that:

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

We now proceed to show that (18) - (20) hold. We start with (18) and consider time instants in $\mathbb{T}^k$. The realized value of i depends on i's stopping time t only when all the other players choose a stopping time above t, that is, when $\theta_{-i}^{\min} > \theta^{*k}(t)$, where $\theta_{-i}^{\min}$ is the smallest signal amongst players other than i:

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

and where $\theta^{*k}(t)$ is the highest type that stops at or before t:

$$\theta^{*k}(t) := \sup\left\{\theta \left| \tau^*\left(h^k, \theta\right) \leq t\right.\right\}.$$

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

$$\frac{\partial V^k(t,\theta)}{\partial t} = \Pr\left(\theta_{-i}^{\min} > \theta^{*k}(t)\right) \cdot \frac{\partial \mathbb{E}\left[v(t,\omega) \left| h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta\right.\right]}{\partial t}. \tag{21}$$

Since $t \in \mathbb{T}^k$, t is the optimal stopping time for type $\theta^{*k}(t)$, i.e. $t = \tau^*\left(h^k, \theta^{*k}(t)\right)$. This means that t is an interior solution to (3), and it follows from optimality that:

$$\frac{\partial \mathbb{E}\left[v(t,\omega) \left| h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta^{*k}(t)\right.\right]}{\partial t} = 0. \tag{22}$$

Supermodularity or log-supermodularity of $v(t,\omega)$ together with MLRP then imply that

$$\frac{\partial \mathbb{E}\left[v(t,\omega) \left| h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta\right.\right]}{\partial t} \geq (\leq)\, 0 \text{ for } \theta > (<)\, \theta^{*k}(t). \tag{23}$$

If $t < (>)\, \tau^*\left(h^k, \theta\right)$, we have $\theta > (<)\, \theta^{*k}(t)$. Therefore (21) and (23) together prove (18).

Next, consider the case where $t \in \left(t^k, \infty\right)$, $t \notin \mathbb{T}^k$. This means that t lies within an interval $t \in (\tau^-, \tau^+)$ such that no signal type invests within it:

$$\begin{aligned}
\tau^- &:= \lim_{\theta' \uparrow \theta^{*k}(t)} \tau^{*k}\left(h^k, \theta'\right) = \sup\left\{t' < t \left| t' \in \mathbb{T}^k\right.\right\}, \\
\tau^+ &:= \lim_{\theta' \downarrow \theta^{*k}(t)} \tau^{*k}\left(h^k, \theta'\right) = \inf\left\{t' > t \left| t' \in \mathbb{T}^k\right.\right\}.
\end{aligned}$$

Notice that signal $\theta^{*k}(t)$ is now a cutoff type such that a signal just above it prefers stopping at τ^+ to t, while a signal type just below it prefers stopping at τ^- to t. Since the information about the other players' signals is summarized by $\theta_{-i}^{\min} > \theta^{*k}(t)$ within the whole interval (τ^-, τ^+), this means that:

$$\begin{aligned}
\lim_{\theta' \downarrow \theta^{*k}(t)} \mathbb{E}\left[v\left(\tau^+, \omega\right) \left| h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta'\right.\right] &\geq \lim_{\theta' \downarrow \theta^{*k}(t)} \mathbb{E}\left[v(t,\omega) \left| h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta'\right.\right], \\
\lim_{\theta' \uparrow \theta^{*k}(t)} \mathbb{E}\left[v\left(\tau^-, \omega\right) \left| h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta'\right.\right] &\geq \lim_{\theta' \uparrow \theta^{*k}(t)} \mathbb{E}\left[v(t,\omega) \left| h^k, \theta_{-i}^{\min} > \theta^{*k}(t), \theta_i = \theta'\right.\right].
\end{aligned}$$

These two equations, supermodularity (or log-supermodularity) of $v(t,\omega)$, MLRP, and the fact that the probability with which some $j \neq i$ stops within (τ^-, τ^+) is zero, imply:

$$\begin{aligned}
V^k\left(\tau^+, \theta\right) &\geq V^k(t,\theta) \text{ if } \theta > \theta^{*k}(t), \text{ and} \\
V^k\left(\tau^-, \theta\right) &\geq V^k(t,\theta) \text{ if } \theta < \theta^{*k}(t).
\end{aligned}$$

To confirm that equations (19) and (20) hold, it only remains to note that if $t > (<)\, \tau^*\left(h^k, \theta\right)$, then $\theta < (>)\, \theta^{*k}(t)$.

Step 2: We argue next that $\tau^*\left(h^k,\theta\right)$ remains optimal in stage k when i plays $\tau^*\left(h^{k+1},\theta\right)$ in stage $k+1$. To see this, note that relaxing the constraint $\tau = t^k$ 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 $\tau^*\left(h^k,\theta\right)$ is preferred to all $t < \tau^*\left(h^k,\theta\right)$ continues to hold when we let i play $\tau^*\left(h^{k+1},\theta\right)$ in stage $k+1$.

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

To summarize: we have now shown that if all players $j \neq i$ play $\tau^*(h,\theta_{-i})$ at all histories h, and if $\tau^*\left(h^{k'},\theta\right)$ is optimal for i in all stages $k' > k$, then $\tau^*\left(h^k,\theta\right)$ is optimal for i in stage k. Since $\tau^*\left(h,\theta\right)$ 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 $\left[Z_1^N, Z_2^N - Z_1^N, ..., Z_k^N - Z_{k-1}^N\right]$ converge to k independent exponential variables for some $k \geq 1$. Consider Z_{k+1}^{N+1}. Since the signals are statistically independent, $\left(\widetilde{\theta}_{k+1}^{N+1} - \widetilde{\theta}_k^{N+1} \middle| \widetilde{\theta}_k^{N+1} = z\right)$ has the same distribution as $\left(\widetilde{\theta}_k^N - \widetilde{\theta}_{k-1}^N \middle| \widetilde{\theta}_{k-1}^N = z\right)$. Multiplying by N we conclude that

$$\left(\frac{N}{(N+1)}(N+1)\left(\widetilde{\theta}_{k+1}^{N+1} - \widetilde{\theta}_k^{N+1}\right) \middle| \widetilde{\theta}_k^{N+1} = z\right)$$

has the same distribution as

$$\left(N\left(\widetilde{\theta}_k^N - \widetilde{\theta}_{k-1}^N\right) \middle| \widetilde{\theta}_{k-1}^N = z\right).$$

Therefore also

$$\left(\frac{N}{(N+1)}\left(Z_{k+1}^{N+1} - Z_k^{N+1}\right) \middle| \widetilde{\theta}_k^{N+1} = z\right)$$

and

$$\left(\left(Z_k^N - Z_{k-1}^N\right) \middle| \widetilde{\theta}_{k-1}^N = z\right)$$

have the same distribution.

By induction hypothesis, $\left(Z_k^N - Z_{k-1}^N\right)$ converges to an exponential random variable, and by the argument above, so does

$$\frac{N}{(N+1)}\left(Z_{k+1}^{N+1} - Z_k^{N+1}\right).$$

Therefore also $\left(Z_{k+1}^N - Z_k^N\right)$ converges to an exponential r.v. as $N \to \infty$. ■

Proof of Lemma 1. Let

$$U(t \mid z_n) := \int_\Omega v\left(t,\omega\right)\pi(\omega \mid z_n)d\omega.$$

Note first that by our assumption of continuity at infinity of $v(t,\omega)$, $t_n(z_n) \subset [0,\infty]$ 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 $U(t \mid z_n) = U(t' \mid z_n)$, we have:

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

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

Proof of Lemma 2. Let

$$U^N(t \mid (z_1,...,z_n)) := \int_\Omega v(t,\omega)\,\pi^N(\omega \mid (z_1,...,z_n))d\omega,$$
$$U^N(t \mid z_n) := \int_\Omega v(t,\omega)\,\pi^N(\omega \mid z_n)d\omega.$$

Consider the sequence $\left\{t_n^N(z_1,...,z_n)\right\}_{N=n}^\infty$. Since

$$\int_\Omega \pi^N(\omega \mid (z_1,...,z_n))d\omega = 1 \text{ for all } N,$$

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

$$\lim_{N\to\infty} U^N(t \mid (z_1,...,z_n)) = U(t \mid z_n). \tag{24}$$

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

$$\lim_{N\to\infty} \frac{\partial U^N(t \mid (z_1,...,z_n))}{\partial t} = \frac{\partial U(t \mid z_n)}{\partial t},$$

and therefore the convergence in equation (24) is uniform. Since $v(t,\omega)$ is continuous at infinity, $U^N(t \mid (z_1,...,z_n))$ has a maximum value. Uniform convergence then implies that

$$\lim_{N\to\infty}\left(\max_t \left(U^N(t \mid (z_1,...,z_n))\right)\right) = \max_t U(t \mid z_n).$$

Take any sequence $\left\{t_n^N\right\}_{N=n}^\infty$ such that $t_n^N \in \arg\max U^N(t \mid (z_1,...,z_n))$ for every N. Since $U^N(t \mid (z_1,...,z_n))$ converges uniformly to $U(t \mid z_n)$, and the latter has a unique maximizer $t_n(z_n)$ by Lemma 1, we have

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

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

Proof of Proposition 3. Fix n and $(z_1,...,z_n)$. Call the player with the i^{th} lowest signal player i. Her normalized signal is z_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 m^{th} lowest, where we must have $m \leq i$. It then follows from Lemma 2 that when $N \to \infty$, the optimal stopping time of i conditional on her information at the time of stopping converges to $t_m(z_i)$, where $m \leq i$. By MLRP and (log)supermodularity of v, we have $t_m(z_i) \geq t_i(z_i)$, and therefore,

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) \geq t_i(z_i). \tag{25}$$

Assume next that

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) > \lim_{N\to\infty} T_{i-1}^N\left(\frac{z_1}{N},...,\frac{z_{i-1}}{N}\right). \tag{26}$$

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

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) = t_i(z_i). \tag{27}$$

Combining equations (25) - (27), we have

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) = \max\left[\lim_{N\to\infty} T_{i-1}^N\left(\frac{z_1}{N},...,\frac{z_{i-1}}{N}\right), t_i(z_i)\right].$$

For the player with the lowest signal, we have:

$$\lim_{N\to\infty} T_1^N\left(\frac{z_1}{N}\right) = t_1(z_1) = \bar{t}_1(z_1).$$

Therefore, it follows by induction that for $i = 2,...,n$

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) = \max\left[\bar{t}_{i-1}(z_1,...,z_{i-1}), t_i(z_i)\right] = \bar{t}_i(z_1,...,z_i).$$

■

Proof of Theorem 2. We analyze the sequence of stopping times $\bar{t}_n(z_1,...,z_n)$, $n = 1,2,...,$ defined by (8) and (9) where the inference is based on exponential random variables. After that, we link those properties to equilibrium stopping times using Corollary 3.

By the weak law of large numbers, the sample average of n exponential random variables Z_1, $Z_2 - Z_1, ..., Z_n - Z_{n-1}$ converges in probability to $1/g(0|\omega)$ as $n \to \infty$. Assumption 3 implies that this identifies the true state ω. Therefore, the unconstrained stopping time $t_n(Z_n)$ defined in (8) converges in probability to the first-best time as $n \to \infty$:

$$t_n(Z_n) \xrightarrow{\mathcal{P}} t^*(\omega). \tag{28}$$

Consider then the distribution of $\bar{t}_n(Z_1,...,Z_n) = \max(t_1(Z_1),...,t_n(Z_n))$. Being the maximum process of t_n, $\bar{t}_n(Z_1,...,Z_n)$ converges in probability to some random variable $\bar{t}_\infty$:

$$\bar{t}_n(Z_1,...,Z_n) \xrightarrow{\mathcal{P}} \bar{t}_\infty, \tag{29}$$

and (28) implies that

$$\lim_{n\to\infty} \Pr\left\{\bar{t}_n(Z_1,...,Z_n) \le t\right\} = 0 \text{ for all } t < t^*(\omega). \tag{30}$$

Consider next the distribution of the first stopping time $t_1(Z_1)$. We have denoted the optimal stopping time under the lowest possible individual signal by $t(0)$. On the other hand, by assumption 3 we have $g(0|\omega) > g(0|\max\Omega)$ for any $\omega < \max\Omega$, and therefore the likelihood ratio across states ω and $\max\Omega$ goes to zero when $z_1 \to \infty$:

$$\lim_{z_1\to\infty} \frac{g(0|\omega)\, e^{-g(0|\omega)z_1}}{g(0|\max\Omega)\, e^{-g(0|\max\Omega)z_1}} = 0.$$

Therefore, we have

$$\lim_{z_1\downarrow 0} t_1(z_1) = t(0) \text{ and } \lim_{z_1\uparrow\infty} t_1(z_1) = \max\Omega,$$

and hence:

$$\lim_{n\to\infty} \Pr\left\{\bar{t}_n(Z_1,...,Z_n) < t(0)\right\} = 0, \tag{31}$$

and

$$\lim_{n\to\infty} \Pr\left\{\bar{t}_n(Z_1,...,Z_n) > t\right\} > 0 \text{ for all } t < \max\Omega. \tag{32}$$

We turn next to the stopping times in the informative equilibrium, and fix a player with signal $\theta > 0$. Consider the game with N players, and let $n(N) = \left\lceil\sqrt{N}\right\rceil$ (where $\lceil\cdot\rceil$ denotes rounding up to the nearest integer). As $N \to \infty$, also $n(N) \to \infty$, so by Corollary 3 and (29), the stopping times of all players that stop after the $n(N)^{th}$ player converge in probability to $\bar{t}_\infty$ as $N \to \infty$. Also, since $n(N)/N \to 0$ as $N \to \infty$, we have

$$\lim_{N\to\infty} \Pr\left\{\widetilde{\theta}^N_{n(N)} < \theta\right\} = 1 \text{ for any } \theta > 0,$$

so that all the players with signals above θ stop later than the $n(N)^{th}$ player. This obviously applies also to the player with the highest signal who stops at time $T^N(\omega)$. Therefore, for any $\theta > 0$,

$$\lim_{N\to\infty} \Pr\{\left|T^N(\omega,\theta) - T^N(\omega)\right| < \varepsilon\} = 1 \text{ for all } \varepsilon > 0,$$

which establishes part 1 of the theorem. Parts 2 and 3 follow then directly from (30), (31), and (32). ■

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.

Dixit, A. & R. Pindyck. 1994. *Investment under Uncertainty*. Princeton: Princeton University Press.

Fudenberg, Drew & Jean Tirole. 1985. "Preemption and Rent Equalization in the Adoption of New Technology." *The Review of Economic Studies* 52(3):pp. 383–401.

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.

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.

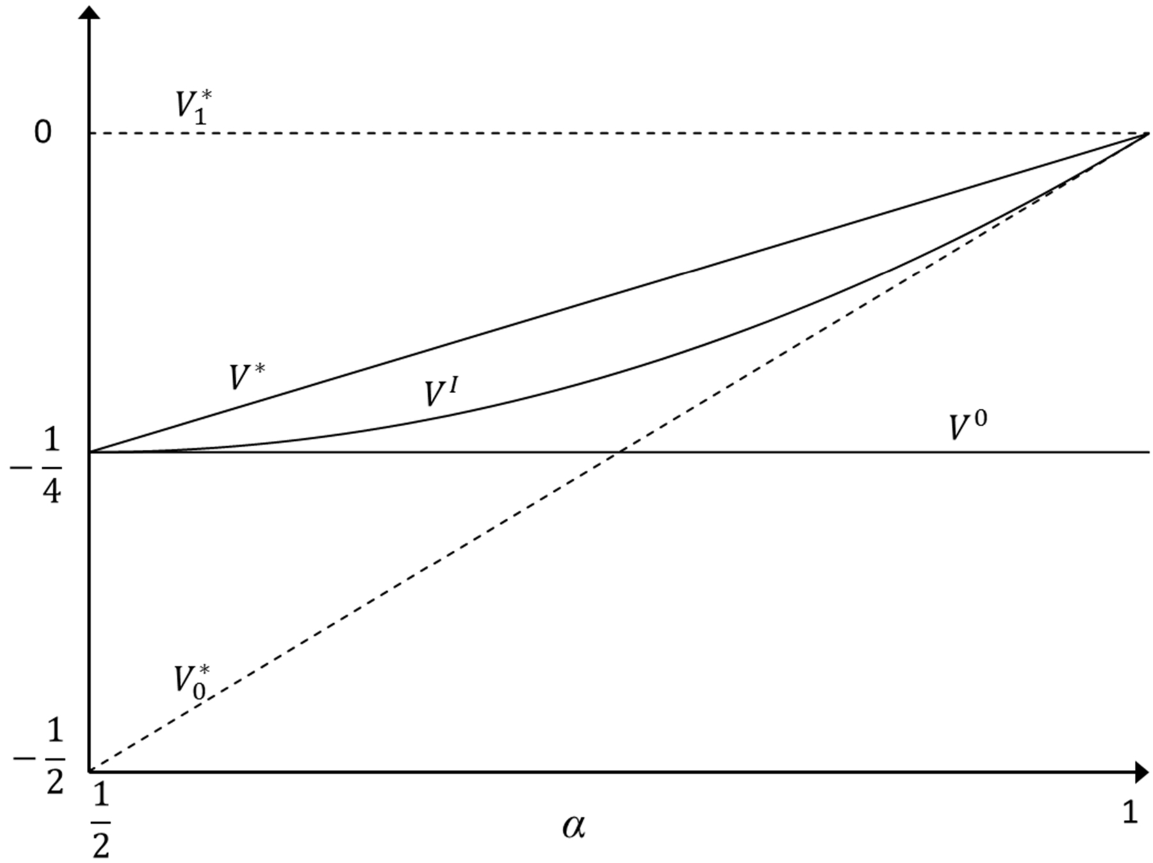

Figure 1: Payoffs as functions of signal precision in the quadratic binary example.

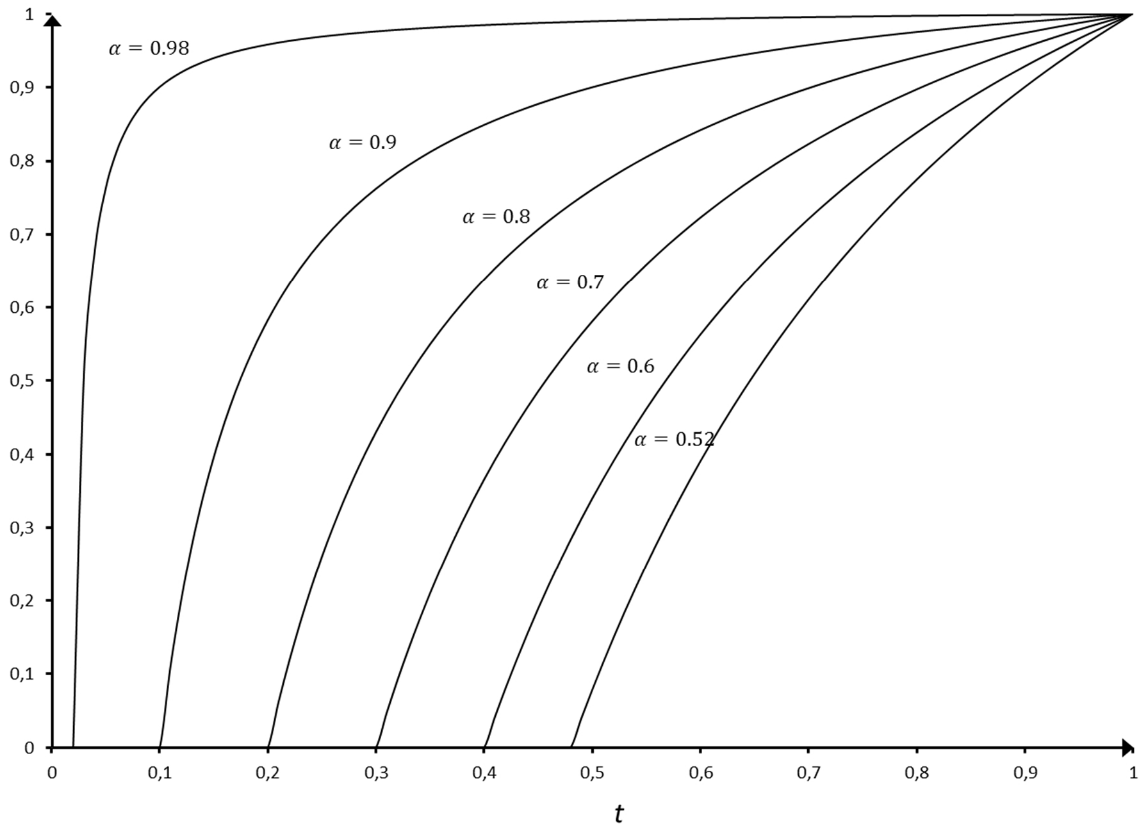

Figure 2: Probability distribution of the stopping time of the last player with various signal precisions ($\omega = 0$).

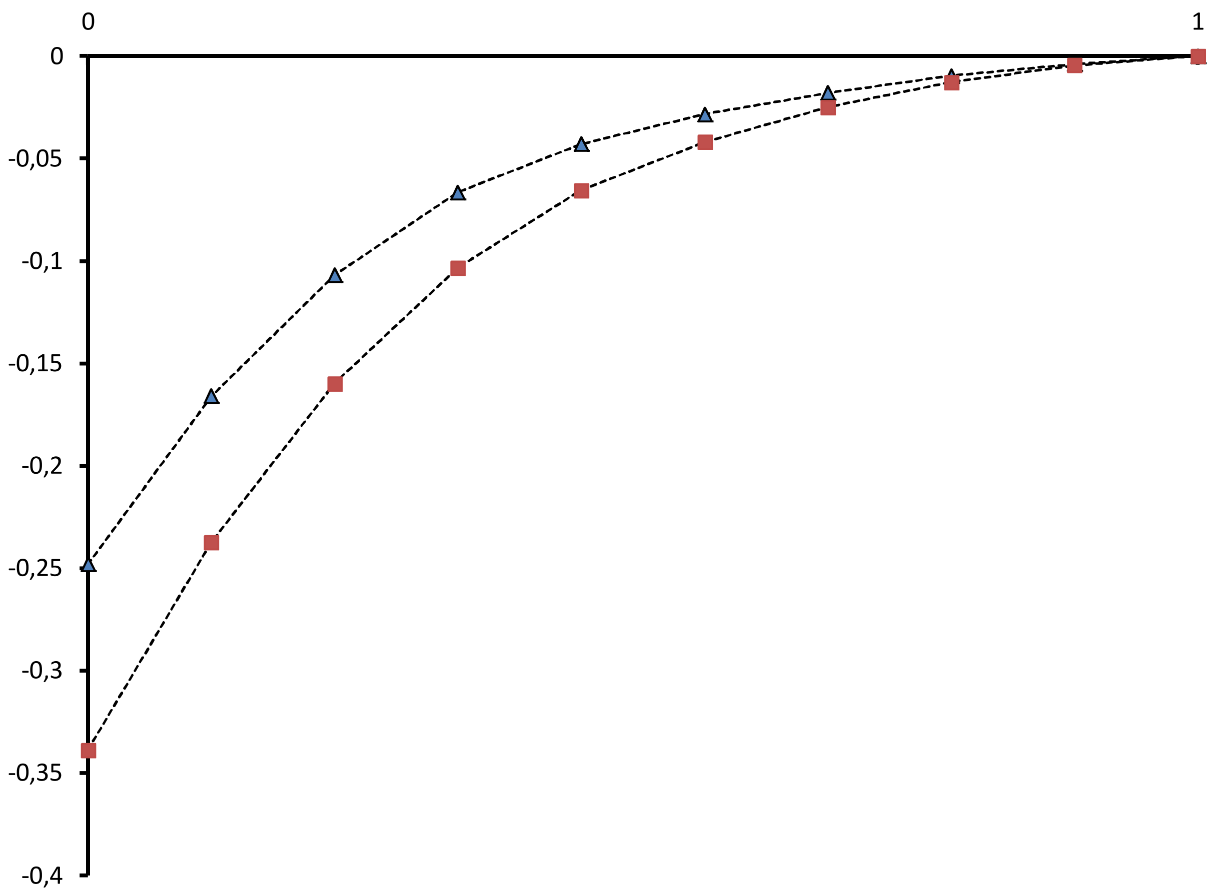

Figure 3: Equilibrium payoffs conditional on state in the ten-state example. Triangle marker: $\alpha = 1$. Square marker: $\alpha = 0.1$.

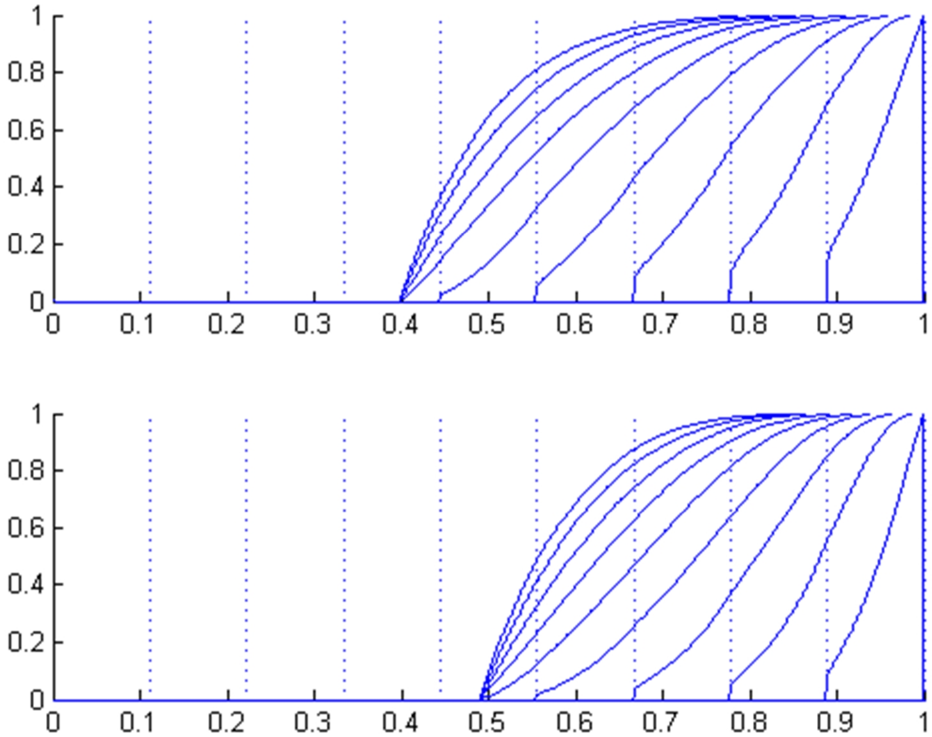

Figure 4: Conditional probability distributions of the stopping time of the last player in the ten-state example. Each solid curve corresponds to one state realization. Dashed lines correspond to first best stopping times for each state. Top panel: $\alpha = 1$. Bottom panel: $\alpha = 0.1$.