Murto, Pauli; Välimäki, Juuso

Delay and information aggregation in stopping games with private information

Published in:
Journal of Economic Theory

DOI:
10.1016/j.jet.2013.07.022

Published: 01/11/2013

Document Version
Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:
Murto, P., & Välimäki, J. (2013). Delay and information aggregation in stopping games with private information. *Journal of Economic Theory*, *148*(6), 2404-2435. https://doi.org/10.1016/j.jet.2013.07.022

Delay and Information Aggregation in Stopping Games with Private Information*

Pauli Murto[†] Juuso Välimäki [‡]

July 24, 2013

Abstract

We consider equilibrium timing decisions in a model with a large number of players and informational externalities. The players have private information about a common payoff parameter that determines the optimal time to invest. They learn from each other in real time by observing past investment decisions. We develop new methods of analysis for such large games, and we give a full characterization of symmetric equilibria. We show that the equilibrium statistical inferences are based on an exponential learning model. Although the beliefs converge to truth, learning takes place too late. Ex-ante welfare is strictly between that without observational learning and that with full information.

1 Introduction

This paper analyzes a large game of timing where the players are privately informed about a common payoff parameter that determines the optimal time to stop the game. The players observe in real time each other's stopping decisions. We assume that the game has a pure informational externality, i.e. the payoff of an individual player does not depend directly on the timing decisions of the other players. We analyze the game in continuous time, and we show how extreme order statistics of the private signals determine equilibrium behavior.

For concreteness, one may interpret the stopping decision as an irreversible investment decision as in the literature on real options. An unknown state variable determines the optimal investment time. The players choose when to invest, and all past investment

*We would like to thank Jacob Goeree, Nicolas Klein, Carlos Oyarzun, Sven Rady, Lones Smith, an associate editor, two referees, and numerous seminar audiences for helpful comments.

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

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

decisions are publicly observable. When deciding whether to invest now or later, a firm takes into account information contained in the other firms' actions. Since the payoff relevant parameter is common to all the players, the equilibrium timing decisions are complementary. Delayed investment by other firms indicates less favorable conditions for early investment whereas early investment by other firms encourages immediate investment. To put it simply, the first firm to invest must always worry about the fact that others have not invested yet. The key question is how the individual players balance the benefits from observing other players' actions with the costs of delay. We analyze how this tension is resolved in the time distribution of the firms' investments.

More specifically, the model is as follows. The first-best time to invest is common to all players and depends on a single state variable ω. Each player has a noisy private signal about ω. The informational setting is 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 log-supermodular in ω and the investment time t. These assumptions ensure a monotonic relationship between a signal and the optimal timing decision based on the signal.

Modeling observational learning in a timing game raises the following issue. When a player invests, the other players' information sets change. As a result, players may want to react immediately to other players' decisions. It is well known that continuous time can be problematic in this respect. An obvious approach is to use discrete time.[1] The drawback of that approach is that it requires complicated limiting procedures if one wants to eliminate the influence of decision lags induced by the discrete periods. In this paper, we avoid this by modeling the dynamic game as a multi-stage game with continuous action sets.

The multi-stage timing game works as follows. At the beginning of each stage, all the remaining players choose their investment 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 leads to 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 other player's investments are captured by allowing investment at time zero of the next stage. While it is well known that for some stopping games with payoff externalities the existence of a stage game equilibrium is problematic in the continuous action variable case, such difficulties do not arise in our game where all externalities are informational.[2]

[1] We adopted that approach in a related paper Murto & Välimäki (2011). We discuss the relationship of that paper to the current one at the end of this section.

[2] An early example of such existence problems appears in Fudenberg & Tirole (1985). With private

We show that the game has symmetric equilibria in monotone strategies. Our characterization result describes a simple method for calculating the optimal decision for each player in the most informative symmetric equilibrium of the game. The key simplifying feature is that while a player takes at each time instant fully into account the information that she has learnt from the past, in equilibrium her decision is not affected by the information that she anticipates to learn in the future. As a result, equilibrium strategy of a player is her optimal stopping time under the assumption that her own signal is the most extreme (that is, favoring early investment) amongst those players that have not yet invested. We show that sometimes the game has also less informative equilibria where all the players invest immediately regardless of their signals.

Our main results concern the limit where the number of players increases towards infinity. A well-known result in Gnedenko (1943) implies that the appropriately normalized k^{th} order statistic in the sample of signals converges in distribution to the sum of independent exponential variables with the parameter determined by the density of the signals at the lower bound of their support. Based on that insight, we show that equilibrium inferences in the limiting model are based on statistical inference on exponentially distributed random variables. This observation remains valid for all large Bayesian games where equilibrium behavior depends on extreme order statistics. We demonstrate this for common value auctions with large numbers of bidders.

Statistical inference in our model is unbiased, but the problem from the players' perspective is that time runs only forward in a dynamic game. Since the players cannot roll back the clock, it is possible that they may have to invest too late relative to their updated beliefs. Obviously the possibility of waiting implies that players do not have to invest too early. This leads to the following results: i) Almost all the players invest too late. ii) Almost all the players invest at the same time. iii) This moment of investment is random conditional on true state. In other words, even though pooled information is perfectly accurate in the large game limit, most of the players end up investing in a big herd that takes place too late and at a random moment.

We show that although part of the potential value of social learning is dissipated through (socially) excessive waiting, the players are ex-ante better off in the informative equilibrium than they would be in isolation from other players. In other words, observational learning improves on average the players' timing decisions. Intuitively, as long as the state remains uncertain, actions are taken by those players that have the most extreme signals. The players with less extreme signals benefit from the information that this generates: they prefer to wait and see how uncertainty resolves before taking any actions themselves.

Related Literature

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.

This paper is closely related to our earlier paper Murto & Välimäki (2011). In that paper, we analyzed a timing game, where players choose when to stop experimenting on bandit processes whose parameters are correlated across players. We already demonstrated in that paper how information is revealed in sudden bursts of activity. Due to the special payoff structure of that game, social learning had no welfare consequences in that model. The current paper shows that once we consider private signals followed by pure observational learning, we can characterize the symmetric equilibria in a much more general environment in terms of information, payoffs, and state space. This identifies the forces behind the qualitative features found in Murto & Välimäki (2011) and at the same time allows us to address welfare consequences of social learning. In the process, we develop technical tools for analyzing the asymptotic properties of equilibrium directly based on extreme order statistics.

More broadly, this paper belongs to a stream of literature that started with models of herding by Banerjee (1992) and Bikhchandani, Hirshleifer & Welch (1992). Those early papers 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 identical, and this leads to qualitatively different pattern of information revelation. Our model has no information cascades, but information is revealed inefficiently late.

The most closely related paper is the investment model by Chamley & Gale (1994).[3] The key difference is that Chamley & Gale (1994) has a payoff function with first-best investment timing either immediately or never, while we allow the state of nature to determine the optimal timing more smoothly, yet capturing Chamley & Gale (1994) as a special case. In other words, they model uncertainty over *whether or not* it is optimal to invest, while we model uncertainty over *when* (if ever) it is optimal to invest. With the payoff structure used in Chamley & Gale (1994), uncertainty is resolved immediately but incompletely at the start of the game whereas our model features gradual information aggregation over time. Section 7 discusses in more detail the relationship between these models.

Moscarini & Squintani (2010) analyze a two-firm R&D race where the inference on common values information is similar to our model. The equilibrium shares with ours the property that the stopping decision by one player may make another player regret that she

[3] 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.

did not stop earlier. Other than that, the results and the analysis in the two papers are quite different. Moscarini & Squintani (2010) focus on the strategic interaction between two players in a model with payoff externalities, whereas our focus is on information aggregation in large games without 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 analyzes the statistical inference problem with a large number of players, and Section 5 analyzes the equilibrium properties in the large game limit. Section 6 presents a quadratic example of the model. Section 7 compares our results to closely related literature and presents some extensions of the basic model. Section 8 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 function

$$v : T \times \Omega \to \mathbb{R}.$$

The state space is a closed subset of the extended real line $\Omega \subseteq [0, \infty]$, and can be either finite or infinite. The players choose their investment time t from the set $T = [0, \infty]$. In order to make sure that optimal choices are well defined, we make the following assumption:

Assumption 1 *The payoff function $v(t, \omega)$ is continuous and bounded on $T \times \Omega$. In particular,*

$$\lim_{t \to \infty} v(t, \omega) = v(\infty, \omega) < \infty \text{ for all } \omega \in \Omega.$$

Furthermore, if $\infty \in \Omega$*, then*

$$\lim_{\omega\to\infty} v(t,\omega) = v(t,\infty) < \infty \text{ for all } t \in T.$$

The players share a common prior $p^0(\omega)$ on Ω and choose the timing of their investment in order to maximize their expectation of v. We assume the following:

Assumption 2 *The payoff function* $v(t,\omega)$ *is piecewise continuously differentiable in* t *for all* ω*. For each* ω*, there is a unique* $t \in [0,\infty]$*, denoted by* $t^*(\omega)$*, 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) Bounded state space[4] and 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, and $0 < r$. 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 $\theta_i \in \Theta = [0,1]$. $G(\theta,\omega)$ 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 3 *For all* i*,* $\theta' > \theta$*, and* $\omega' > \omega$*,*

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

Assumptions 2 and 3 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 $g(0 \mid \omega) \neq g(0 \mid \omega')$

[4] To keep payoff bounded, choose some $\bar{t} > \max\Omega$ and let $v(t,\omega) = v(\bar{t},\omega)$ for $t > \bar{t}$.

[5] A variant of this model with a stochastic state variable will be discussed in Section 7.3.

whenever $\omega \neq \omega'$ (note that assumption 3 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 perfectly 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 4 *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 $\tau_i(h^0, \theta_i) \geq 0$ 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:

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

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. 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

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

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

A pure behavior strategy for stage k is a function

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

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 higher than $\underline{\theta}^k$, 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:

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

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:*

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

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 and Theorem 1 are proved in the Appendix. ■

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

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

In words, $\theta^{*k}\left(t\right)$ 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 $\theta^{*k}\left(t\right)$ (i.e. the lower bound of the signals of the existing players) moves upward as time goes by. As long as no player stops, this implies by MLRP and the (log)supermodularity of v that the optimal stopping time conditional on current information moves forward 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, $\theta^{*k}\left(t\right)$ 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 5 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, all the players stop immediately irrespective of 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.

There are no other types of symmetric equilibria. Note first that any symmetric equilibrium must be monotonic, otherwise a low signal player who stops later than a high signal player would gain by mimicking the high type (or vice versa). Suppose that player i with signal $\theta_i = \theta$ stops at some $t > t^{k-1}$ in a symmetric equilibrium at stage k of the game. Then t must be an optimal stopping time conditional on the information that i has at time t about the other players. By strict (log-)supermodularity, a player with $\theta_j < \theta$ should stop strictly earlier with this information, while a player with $\theta_j > \theta$ should stop strictly later. It follows that in any symmetric equilibrium with $t > t^{k-1}$, i chooses the best stopping time conditional on $\theta_j \geq \theta$ for all $j \neq i$. In other words, any best-response $t > t^{k-1}$ must be as in Theorem 1.[6] The only possible symmetric equilibria are then the uninformative equilibrium, where all the players stop at $t = t^{k-1}$, and the informative equilibrium defined in (3).

It should also 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 7.1.

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. All symmetric equilibria in the game take the form where $\tau^*\left(h^k, \theta\right)$ is played in the first $k-1$ stages, followed by an uninformative equilibrium play in stage k (ending the game). For stage k, the informative equilibrium payoff exceeds the uninformative equilibrium payoff. This follows immediately from the observation that the uninformative equilibrium strategy is also available in the informative equilibrium and gives the same payoff to i regardless of the strategies of players $j \neq i$. It follows that the best symmetric equilibrium is the one where $\tau^*\left(h^k, \theta\right)$ is played in all stages.

4 Statistical inference

In this section, we analyze statistical inference based on the extreme order statistics of the players' signals as we increase the number of players towards infinity. The findings will be utilized in Section 5 where we analyze the informative equilibrium in the large game limit.

To see why order statistics are important for our analysis, note that the informative equilibrium strategy is monotonic in signals, and therefore the players stop in the ascending order of their signals. Hence, the real time instant at which the n:th player stops the game is a function of the n lowest signal realizations amongst the players.

We start with a statistical observation regarding the distribution of extreme order statistics in large samples (Section 4.1). We then show how this translates into optimal stopping times (Section 4.2). Finally, we show how our findings can be utilized to obtain

[6] In Theorem 1, we have chosen the minimum of the best responses as an equilibrium strategy. This choice is inconsequential for the expected payoffs in the game.

new results in other Bayesian games, in particular auctions (Section 4.3).

4.1 Extreme order statistics

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

$$\widetilde{\theta}_i^N := \min\left\{\theta \in [0,1] \mid \#\left\{j \in \{1,...,N\} \mid \theta_j \leq \theta\right\} = i\right\}. \tag{5}$$

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 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 as $N \to \infty$, learning from the order statistics is equivalent to sampling independent random variables from an exponential distribution with an unknown state-dependent parameter $g(0 \mid \omega)$. To get intuition for this result, note that 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 density 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.[7] Bayes' rule and simple algebra show that these posteriors are equal:

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

$$= \frac{\pi^0(\omega) \cdot (g(0 \mid \omega))^n e^{-g(0|\omega) z_n}}{\int_\Omega \pi^0(\omega') \cdot (g(0 \mid \omega'))^n e^{-g(0|\omega') z_n} d\omega'} = \pi(\omega \mid z_n).$$

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 *Fix a sample of normalized order statistics* $(z_1, ..., z_n)$. *Then*

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

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 implies 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 will play 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 1, she would now choose the optimal stopping time conditional on her information on those

[7] The case with a discrete Ω is handled similarly.

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 Optimal timing based on order statistics

We now consider optimal stopping times based on the inference on order statistics. This is a purely hypothetical problem, where the decision maker is unconstrained in the sense that she can choose any stopping time in $[0, \infty]$. This is in contrast to the actual timing game defined in Section 2.2, where a player in stage k is constrained to choose a stopping time in $\left[t^{k-1}, \infty\right]$. We link the results of this section to the timing game in Section 5.

First, consider inference based on 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 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 proposition, 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.

Proposition 3 *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. ■

4.3 Other applications

It should be noted that the statistical inference derived in Section 4.1 is valid for many other large Bayesian games. Consider, for example, a $k+1^{st}$ price auction of k identical objects as specified in Pesendorfer & Swinkels (1997). There are N bidders with unit demands, $\omega \in [0,1]$ denotes the common value of the object for sale, and $\theta_i \in [0,1]$ is the signal of bidder i. To keep notation consistent with the rest of the paper, we adopt a non-standard notation where a low signal value indicates high value for the object.[8] Pesendorfer & Swinkels (1997) have shown that in the unique equilibrium of such a game, the equilibrium bid $b^N(\theta)$ is the expected value of the object conditional on the event that $\widetilde{\theta}_k^{N-1} = \theta_i = \theta$:

$$b^N(\theta) = \mathbb{E}[\omega \left| \widetilde{\theta}_k^{N-1} = \theta_i = \theta \right.].$$

Using Proposition 2 and its corollaries, we note that in the limit $N \to \infty$, the posterior on ω conditional on $N \cdot \widetilde{\theta}_k^{N-1} = N \cdot \theta_i = z$ converges to the posterior on ω conditional on $Z^{k+1} = z$, where $Z^{k+1} \sim \Gamma(k+1, g(0 \mid \omega))$:

$$\begin{aligned}\lim_{N\to\infty} \pi^N(\omega \quad | \quad N \cdot \theta_i = z, N \cdot \widetilde{\theta}_k^{N-1} = z) &= \pi(\omega \mid z) \\ &= \frac{\pi^0(\omega) \cdot (g(0 \mid \omega))^{k+1} e^{-g(0|\omega)z_{k+1}}}{\int_0^1 \pi^0(\omega') \cdot (g(0 \mid \omega'))^{k+1} e^{-g(0|\omega')z_{k+1}} d\omega'}.\end{aligned}$$

Therefore, the bid function for normalized signal $N\theta = z$ converges to

$$\lim_{N\to\infty} b^N(z) := b(z) = \frac{\int_0^1 \omega' \pi^0(\omega') \cdot (g(0 \mid \omega'))^{k+1} e^{-g(0|\omega')z} d\omega'}{\int_0^1 \pi^0(\omega') \cdot (g(0 \mid \omega'))^{k+1} e^{-g(0|\omega')z} d\omega'}.$$

The realized price is then a random variable $P = b(z)$, where $z \sim \Gamma(k+1, g(0 \mid \omega))$. To our knowledge, this limit price distribution has not appeared in the literature before.

Similar explicit calculations can be performed for other auction formats including first price (or descending) auction. Since the equilibrium statistical inference is quite simple, it is also possible to extend the analysis beyond the risk neutral case. In Murto & Välimäki (2012), we compare the expected revenues for different auction formats under CARA preferences.

5 Informative equilibrium in large games

In this section, we relate the optimal stopping times derived in Section 4.2 to the equilibrium stopping times. Since the informative equilibrium strategy defined in Theorem 1 is monotonic in signals, the players stop in the ascending order of their signals. Therefore,

[8] Translated to the notation used in Pesendorfer & Swinkels (1997), our θ corresponds to $1-s$ in their paper. Other than that, we assume here that the model is exactly as in Pesendorfer & Swinkels (1997).

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 $\widetilde{\theta}_i^N$ is the i^{th} order statistic in the game with N players. We let $N \to \infty$, and derive the statistical properties of $T_n^N\left(\widetilde{\theta}_1^N, ..., \widetilde{\theta}_n^N\right)$ for small n in Section 5.1 below, and for large n in Section 5.2 below. We discuss the welfare properties of the informative equilibrium in Section 5.3.

5.1 First n stopping times

In Section 4.2, we defined the optimal stopping time $t_n(z_n)$ of the limit model based on the normalized n^{th} order statistic. Since the sequence $\{t_n(z_n)\}$ is not necessarily increasing in n, we call this the *unconstrained* stopping time. In the actual timing game the players cannot go backwards in time and a more relevant stopping time for the player with the n^{th} lowest signal is the *constrained* stopping time,

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

where $t_{n'}(z_{n'})$ is defined in (8). If the constrained stopping time for the n^{th} order statistic is different from the unconstrained one, then the player with the n^{th} lowest signal stops immediately at the beginning of some stage and regrets her decision not to stop already at an earlier stage.

The main result of this section is that the stopping times in the informative equilibrium of the game converge to the constrained stopping times defined in (9). We have:

Proposition 4 *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. ■

As a corollary to this result, we can relate 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: 1) fix an arbitrary n, 2) draw n independent realizations $(z_1, ..., z_n)$ from exponential distribution with parameter $g(0 \mid \omega)$, and 3) 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 4. ■

5.2 Delay in equilibrium

We now 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. The randomness in the stopping time reflects the stopping decisions of the other players in the game. We are 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 limit is to be understood in the sense of convergence in distribution.

The key random variable for understanding the informative equilibrium is the time instant at which the *last* player stops, denoted by $T^N(\omega)$. Again, we consider the large game limit:

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

We seek a characterization of $T(\omega)$, and furthermore, we argue that the ex ante expected payoff to the players in state ω converges to the expectation of $v(T(\omega), \omega)$ as N grows.

We let $F(t \mid \omega)$ denote the distribution of $T(\omega)$:

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

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

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

Proof. In the Appendix. ■

Theorem 2 summarizes 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 statistical 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.

5.3 Welfare

Denote by V^* the ex-ante value of a player in the informative equilibrium of the game. To address welfare implications of observational learning, we contrast V^* to the ex-ante value of an isolated player that chooses optimal stopping time based on her own signal only.

Denote by $t(\theta)$ the optimal stopping time of an isolated player with signal θ:

$$t(\theta) := \arg\max_{t\in[0,\infty]} \mathbb{E}_\omega \left[v(t,\omega) \,|\theta_i = \theta\right].$$

The ex-ante value of an isolated player is then:

$$V^I := \mathbb{E}_\theta \left[\mathbb{E}_\omega \left(v(t(\theta),\omega) \,|\theta_i = \theta\right)\right].$$

Consider a player with signal θ such that $t(\theta) > t(0)$. One feasible (non-optimal) strategy for this player is to stop at time $t(\theta)$ if at least one other player still remains in the game at that moment, or stop immediately after the last player other than her has stopped. The realized stopping time resulting from this strategy is

$$\widetilde{t}(\theta) = \min\left(T(\omega), t(\theta)\right).$$

From part 2 of the Theorem 2 we see that if $t^*(\omega) < t(\theta)$, then with positive probability $t^*(\omega) < T(\omega) < t(\theta)$. Therefore, stopping at $\widetilde{t}(\theta)$ is strictly better than stopping at $t(\theta)$ for some state realizations. Since by the same Theorem, $\Pr\left(T(\omega) < t^*(\omega)\right) = 0$, stopping at $\widetilde{t}(\theta)$ can never be worse than stopping at $t(\theta)$. Since the equilibrium strategy cannot be worse than this non-optimal strategy, it follows that $V^* > V^I$. This reasoning formalizes the intuition that a player with a high signals benefits from the information revealed by the actions of those players that have lower signals.

On the other hand, we also see from Theorem 2 that $\Pr(T(\omega) > t^*(\omega)) > 0$, so we must have $V^* < \overline{V}$, where $\overline{V}$ is the ex-ante payoff with perfect information on state:

$$\overline{V} := \mathbb{E}_\omega \left(v \left(t^* \left(\omega \right), \omega \right) \right).$$

We have thus derived loose bounds for the ex-ante payoff in equilibrium:

$$V^I < V^* < \overline{V}.$$

Unfortunately, we are not able to say more in general. The difficulty is that while our supermodularity assumption restricts the direction to which the maximizer of v changes as information changes, it puts no restriction on how the expected *level* of v changes.

In order to conduct a more complete welfare analysis, we must specify the model further. In the next section we show that by restricting to a binary-state case (as most literature on social learning), we are able to pin down equilibrium payoffs conditional on state. Combined with a quadratic payoff function, we get analytic expressions for equilibrium payoffs allowing natural comparative statics. With more than two states, we can compute the payoffs by simulation.

6 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 there are only two possible states. This allows us to compute analytically the welfare implications of observational learning. For simplicity, we also assume signals to be essentially binary, although this is not important for the analysis.

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). We call signals below (above) θ^* low (high) and write $\theta = l (= h)$. 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, this is equivalent to a standard binary-signal, binary-state model, where α measures the precision of the signals. We assume that the prior probability is $p^0 = \Pr\{\omega = 1\} = \frac{1}{2}$.

The payoffs are given by[9]

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

Hence the optimal action for a player with posterior p on $\{\omega = 1\}$ is to invest at $t = p$. The ex-ante payoff with perfect information is $\overline{V} = 0$.

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 private signal. Then she stops at $t = 1/2$ and her payoff is

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

Second, suppose that an isolated player observes her own private signal. 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 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. $V^I \uparrow \overline{V} = 0$ 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 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).

[9] To keep payoff bounded from below, let $v(t, \omega) = -(1 - \omega)^2$ for all $t > 1$.

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 $T(\omega)$, 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 in accordance with Section 5.3, $V^I < V^* < \overline{V} = 0$ whenever $\alpha \in (\frac{1}{2}, 1)$. 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 ($V^* < \overline{V}$). 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 $T(\omega)$. As long as $t > 1-\alpha$ and some of the uninformed players with a low signal 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 $p_{\theta=l}(t) = 0$ 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 $T(\omega)$, 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. We see that the more precise the signals, the higher the hazard rates.

< 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 5.1, 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 the 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.

7 Discussion

7.1 Relation to Chamley and Gale (1994)

Our paper extends the models in Chamley & Gale (1994) and Chamley (2004) to a more general payoff specification. To understand the relationship between the models, 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

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

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.

7.2 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.

7.3 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.

8 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.

9 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 then consider an arbitrary history h^k and assume that player i deviates from $\tau^*\left(h^k,\theta\right)$ to an arbitrary $t \geq t^{k-1}$ in stage k, but uses $\tau^*\left(h^{k'},\theta\right)$ for all $k' > k$. We denote by $\widetilde{V}\left(t\right)$ the value of this deviation evaluated at the beginning of stage k. Our goal is to show that $\tau^*\left(h^k,\theta\right) \in \arg\max_{t\geq t^{k-1}}\widetilde{V}\left(t\right)$.

Let $\theta_{-i}^{\min} := \min_{j\in\mathcal{N}^k\setminus i}\theta_j$ denote the smallest signal amongst players other than i and recall from (4) that $\theta^{*k}(t)$ denotes the highest type that stops at or before time t. We may now formally express the value from deviation. For $t = t^{k-1}$, we have:

$$\widetilde{V}\left(t^{k-1}\right) = \mathbb{E}\left[v\left(t^{k-1},\omega\right)\middle|h^k,\theta_i=\theta\right], \tag{18}$$

and for $t > t^{k-1}$:

$$\begin{aligned}\widetilde{V}(t) \;=\;& \Pr\left(\theta_{-i}^{\min} > \theta^{*k}(t)\middle|h^k,\theta_i=\theta\right)\mathbb{E}\left[v(t,\omega)\middle|h^k,\theta_i=\theta,\theta_{-i}^{\min}>\theta^{*k}(t)\right] \\ &+\Pr\left(\theta_{-i}^{\min}\le\theta^{*k}(t)\middle|h^k,\theta_i=\theta\right)\mathbb{E}\left[V\left(h^{k+1}\right)\middle|h^k,\theta_i=\theta,\theta_{-i}^{\min}\le\theta^{*k}(t)\right],\end{aligned} \tag{19}$$

where $V\left(h^{k+1}\right)$ denotes the continuation value of i at the beginning of stage $k+1$, when following $\tau^*\left(h^{k'},\theta\right)$ for all stages $k' > k$.

It is useful to set an upper bound for potential profitable deviations. Let $\bar{t}$ denote the optimal stopping time assuming that all the players have the highest possible signal:

$$\bar{t} := \min\left(\arg\max_{t\ge t^{k-1}}\mathbb{E}\left[v(t,\omega)\middle|h^k,\theta_j=1\text{ for all }j\in\mathcal{N}^k\right]\right).$$

We organize the proof in four claims. Claim 1 establishes the continuity of $\widetilde{V}(t)$ in $(t^{k-1},\bar{t}]$. Claim 2 together with Claim 1 shows that $\arg\max_t \widetilde{V}(t) \ge \tau^*\left(h^k,\theta\right)$. Claims 3 and 4 together with Claim 1 show that $\min\left(\arg\max_t \widetilde{V}(t)\right) \le \tau^*\left(h^k,\theta\right)$.

Claim 1: $\widetilde{V}(t)$ is continuous and bounded within $\left(t^{k-1},\bar{t}\right]$.

Proof of Claim 1: Whenever $\tau^*\left(h^k,\theta\right) \in \left(t^{k-1},\bar{t}\right)$, then by MLRP and strict (log)supermodularity of v, $\tau^*\left(h^k,\theta\right)$ is strictly increasing in θ. It follows that $\theta^{*k}(t)$ is continuous within $\left(t^{k-1},\bar{t}\right)$. Since $v(t,\omega)$ is continuous in t for all ω, and each of the terms in (19) is continuous in $\theta^{*k}(t)$, it follows that $\widetilde{V}(t)$ is continuous in t within $\left(t^{k-1},\bar{t}\right)$. Although it is possible that $\theta^{*k}(t)$ is discontinuous at $t=\bar{t}$,[10] we must have $\tau^*\left(h^k,\theta\right)\le\bar{t}$ for all θ, and hence $\theta^{*k}\left(\bar{t}\right)=1$. This is because $\bar{t}$ is an upper bound for any optimal stopping time in our model. Since all the remaining players stop at latest at $\bar{t}$ irrespective of their signals, it is not possible to observe any new information at that moment. Continuity of $\widetilde{V}(t)$ at $\bar{t}$ then follows directly from continuity of v. $\widetilde{V}(t)$ is trivially bounded because by assumption v is bounded.

Claim 2:

$$\text{For all } t\in\left[t^{k-1},\tau^*\left(h^k,\theta\right)\right),\text{ there is some } t'\in\left(t,\tau^*\left(h^k,\theta\right)\right] \text{ such that } \widetilde{V}(t') > \widetilde{V}(t).$$

Proof of Claim 2: Denote by $t^*(\theta,\theta')$ the optimal stopping time of a player with signal θ, when all the other players are known to have signals above θ':

$$t^*(\theta,\theta') := \min\left(\arg\max_{t\ge t^{k-1}}\mathbb{E}\left[v(t,\omega)\middle|h^k,\theta_i=\theta,\theta_{-i}^{\min}\ge\theta'\right].\right) \tag{20}$$

[10] This happens if signals become uninformative above some threshold $\theta^* < 1$ as in our binary signal example of Section 6.

If $t \in \left[t^{k-1}, \tau^*\left(h^k, \theta\right)\right)$, then $\theta > \theta^{*k}(t)$, and by MLRP and strict (log-)supermodularity of v, we have $t < t^*\left(\theta, \theta^{*k}(t)\right) < \tau^*\left(h^k, \theta\right)$. In other words, conditional on information available at time t, it is strictly better for i to stop at $t^*\left(\theta, \theta^{*k}(t)\right)$ than t. Any possibility to revise this decision in the interval $\left(t, t^*\left(\theta, \theta^{*k}(t)\right)\right)$ may only reinforce this conclusion. It follows that $\widetilde{V}\left(t^*\left(\theta, \theta^{*k}(t)\right)\right) > \widetilde{V}(t)$.

Claim 3:

$$\text{For all } t \in \left(\tau^*\left(h^k, \theta\right), \bar{t}\right], \text{ there is some } t' \in \left[\tau^*\left(h^k, \theta\right), t\right) \text{ such that } \widetilde{V}(t') \geq \widetilde{V}(t).$$

Proof of Claim 3: Since $t > \tau^*\left(h^k, \theta\right)$, we have $\theta \leq \theta^{*k}(t)$. Suppose first that this holds as equality: $\theta = \theta^{*k}(t)$. Then $\theta^{*k}(t') = \theta$ for all $t' \in \left[\tau^*\left(h^k, \theta\right), t\right]$. In that case i learns nothing from others within $\left[\tau^*\left(h^k, \theta\right), t\right]$ since no player stops there. By definition, $\tau^*\left(h^k, \theta\right)$ is an optimal stopping time for a player with that information, therefore $\widetilde{V}\left(\tau^*\left(h^k, \theta\right)\right) \geq \widetilde{V}(t)$.

Consider next strict inequality $\theta < \theta^{*k}(t)$. Let $t' = t^*\left(\theta, \theta^{*k}(t)\right) < t$ and let $\theta' = \theta^{*k}(t')$. Here t' represents the optimal stopping time of i conditional on her information at time t, and θ' represents the cutoff type that stops at time t in equilibrium. We have $\theta < \theta' < \theta^{*k}(t)$ and $\tau^*\left(h^k, \theta\right) < t' < t$. We want to show that stopping at t' dominates stopping at t. To do that, we decompose (19) according to whether $\theta_{-i}^{\min}$ lies in $(0, \theta')$, $\left(\theta', \theta^{*k}(t)\right)$ or $\left(\theta^{*k}, 1\right)$, and compare $\widetilde{V}(t)$ and $\widetilde{V}(t')$ term by term.

Suppose that $\theta_{-i}^{\min} \in \left(\theta', \theta^{*k}(t)\right)$. Then $t' < \tau^*\left(h^k, \theta_{-i}^{\min}\right) < t$, so that choosing t' means that i is the first to stop, while choosing t means that another player j stops first at time $\tau^*\left(h^k, \theta_{-i}^{\min}\right) < t$. Since $\tau^*\left(h^k, \theta_{-i}^{\min}\right) > \tau^*\left(h^k, \theta\right)$, Proposition 1 implies that $\tau^*\left(h^{k+1}, \theta\right) = t^k$, so that i will stop immediately at the beginning of the next stage and the continuation value $V\left(h^{k+1}\right)$ is the expectation of $v\left(\tau^*\left(h^k, \theta_{-i}^{\min}\right), \omega\right)$. Therefore, evaluating (19) at t gives

$$\begin{aligned}
\widetilde{V}(t) = & \Pr\left(\theta_{-i}^{\min} > \theta^{*k}(t)\right) \mathbb{E}\left[v(t, \omega) \left|\theta_{-i}^{\min} > \theta^{*k}(t)\right.\right] \\
& + \Pr\left(\theta' < \theta_{-i}^{\min} \leq \theta^{*k}(t)\right) \mathbb{E}\left[v\left(\tau^*\left(h^k, \theta_{-i}^{\min}\right), \omega\right) \left|\theta' < \theta_{-i}^{\min} \leq \theta^{*k}(t)\right.\right] \\
& + \Pr\left(\theta_{-i}^{\min} \leq \theta'\right) \mathbb{E}\left[V\left(h^{k+1}\right) \left|\theta_{-i}^{\min} \leq \theta'\right.\right],
\end{aligned}$$

while evaluating (19) at t' gives

$$\begin{aligned}
\widetilde{V}(t') = & \Pr\left(\theta_{-i}^{\min} > \theta^{*k}(t)\right) \mathbb{E}\left[v(t', \omega) \left|\theta_{-i}^{\min} > \theta^{*k}(t)\right.\right] \\
& + \Pr\left(\theta' < \theta_{-i}^{\min} \leq \theta^{*k}(t)\right) \mathbb{E}\left[v(t', \omega) \left|\theta' < \theta_{-i}^{\min} \leq \theta^{*k}(t)\right.\right] \\
& + \Pr\left(\theta_{-i}^{\min} \leq \theta'\right) \mathbb{E}\left[V\left(h^{k+1}\right) \left|\theta_{-i}^{\min} \leq \theta'\right.\right],
\end{aligned}$$

where Pr and $\mathbb{E}$ should here be understood conditional on h^k and $\theta_i = \theta$. By strict (log-)supermodularity and MLRP, we have

$$\begin{aligned}
\mathbb{E}\left[v(t', \omega) \left|\theta_{-i}^{\min} > \theta^{*k}(t)\right.\right] & > \mathbb{E}\left[v(t, \omega) \left|\theta_{-i}^{\min} > \theta^{*k}(t)\right.\right] \text{ and} \\
\mathbb{E}\left[v(t', \omega) \left|\theta' < \theta_{-i}^{\min} \leq \theta^{*k}(t)\right.\right] & > \mathbb{E}\left[v\left(\tau^*\left(h^k, \theta_{-i}^{\min}\right), \omega\right) \left|\theta' < \theta_{-i}^{\min} \leq \theta^{*k}(t)\right.\right].
\end{aligned}$$

Therefore, $\widetilde{V}(t') > \widetilde{V}(t)$.

Claim 4:

$$\widetilde{V}(t) \leq \widetilde{V}(\bar{t}) \text{ for all } t > \bar{t}.$$

Proof of Claim 4: $\bar{t}$ is the optimal stopping time conditional on all the players having the highest possible signal. By strict (log-)supermodularity and MLRP, stopping at $\bar{t}$ is preferable to stopping at $t \geq \bar{t}$ conditional on any information that i might learn from the other players, therefore $\widetilde{V}(t) \leq \widetilde{V}(\bar{t})$ for all $t > \bar{t}$.

We may now finish the proof by combining the above results. Claim 2 implies that $\tau^*(h^k, \theta)$ is the maximizer of $\widetilde{V}(t)$ on any interval $[t, \tau^*(h^k, \theta)]$ for $t > t^{k-1}$. By Claim 1, this property extends to $[t^{k-1}, \tau^*(h^k, \theta)]$:

$$\tau^*(h^k, \theta) \in \arg \max_{t \in [t^{k-1}, \tau^*(h^k, \theta)]} \widetilde{V}(t).$$

Claims 1 and 3 imply that

$$\tau^*(h^k, \theta) \in \arg \max_{t \in [\tau^*(h^k, \theta), \bar{t}]} \widetilde{V}(t).$$

Combining these with Claim 4 gives

$$\tau^*(h^k, \theta) \in \arg \max_{t \geq t^{k-1}} \widetilde{V}(t).$$

We have now shown that if all players $j \neq i$ play $\tau^*(h, \theta_{-i})$ at all histories h, and if $\tau^*(h^{k'}, \theta)$ is optimal for i in all stages $k' > k$, then $\tau^*(h^k, \theta)$ is optimal for i in stage k. Since $\tau^*(h, \theta)$ 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 $[Z_1^N, Z_2^N - Z_1^N, ..., Z_k^N - Z_{k-1}^N]$ 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

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

Since $v\left(t,\omega\right)$ is continuous on $T\times\Omega$, we note that also $U(t \mid z_n)$ is continuous in t (including continuity at $t=0$ and $t=\infty$). Therefore, a maximizer exists and $t_n\left(z_n\right) \subset [0,\infty]$ is non-empty. For the uniqueness, we use Theorem 1 in Araujo & Mas-Colell (1978). To this effect, we note from (7) that $\pi(\omega \mid z_n)$ is continuously differentiable in z_n. Using the functional form of $\frac{\partial}{\partial z_n}\pi\left(\omega \mid z_n\right)$ computed from (7), and noting that $v\left(t,\omega\right)$ and $g\left(0 \mid \omega\right)$ are bounded, it is easy to find a constant M that guarantees

$$\left|\frac{\partial}{\partial z_n}\left[v\left(t,\omega\right)\pi\left(\omega \mid z_n\right)\right]\right| < M\pi\left(\omega \mid z_n\right)$$

for all $\omega \in \Omega$ and for all $z_n \geq 0$. Since $\pi\left(\omega \mid z_n\right)$ is a probability distribution, $M\pi\left(\cdot \mid z_n\right): \Omega \to \mathbb{R}$ is integrable, and therefore the derivative of $U(t \mid z_n)$ with respect to z_n exists and can be obtained by differentiating under the integral sign (e.g. Lemma 2.2 in Lang (1983)):

$$\frac{\partial}{\partial z_n}U(t \mid z_n) = \int_\Omega v\left(t,\omega\right)\frac{\partial}{\partial z_n}\pi\left(\omega \mid z_n\right)d\omega.$$

Since $v\left(t,\omega\right)$ is continuous and $\pi(\omega \mid z_n)$ is continuously differentiable in z_n, $\frac{\partial}{\partial z_n}U(t \mid z_n)$ is continuous both in t and z_n. Furthermore, MLRP and the strict (log-)supermodularity of $v\left(t,\omega\right)$ 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 Proposition 3. We prove here that for almost every $(z_1,...,z_n)$, $\lim_{N\to\infty} t_n^N\left(z_1,...,z_n\right) = t_n\left(z_n\right)$. The proof is identical for $\lim_{N\to\infty} t_n^N\left(z_n\right) = t_n\left(z_n\right)$.

Let

$$U^N(t \mid (z_1,...,z_n)) := \int_\Omega v\left(t,\omega\right)\pi^N(\omega \mid (z_1,...,z_n))d\omega \text{ and}$$

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

Since $v(t,\omega)$ is continuous on $T\times\Omega$, both $U^N(t \mid (z_1,...,z_n))$ and $U(t \mid z_n)$ are continuous in t and have maximizers in $T=[0,\infty]$. Let K denote a bound of $|v(t,\omega)|$ on $T\times\Omega$. We start by showing that $U^N(t \mid (z_1,...,z_n))$ converges uniformly to $U(t \mid z_n)$:

$$\begin{aligned} &\lim_{N\to\infty}\sup_{t\in T}\left|\int_\Omega v(t,\omega)\,\pi^N(\omega \mid (z_1,...,z_n))d\omega - \int_\Omega v(t,\omega)\,\pi(\omega\,|z_n)\,d\omega\right| \\ \leq\;& K\lim_{N\to\infty}\int_\Omega \left|\pi^N(\omega \mid (z_1,...,z_n)) - \pi(\omega\,|z_n)\right| d\omega = 0, \end{aligned}$$

where the last equality follows from Scheffé's Theorem (e.g. Theorem 16.11 in Billingsley (1986)) and the fact that $\pi^N(\omega \mid (z_1,...,z_n))$ converges pointwise to $\pi(\omega\,|z_n)$.

Take any sequence $\{t^N(z_1,...,z_n)\}_{N=n}^{\infty}$ such that $t^N(z_1,...,z_n) \in \arg\max U^N(t \mid (z_1,...,z_n))$ for every N. It is easy to see that for all N, we have

$$\left|U^N(t^N(z_1,...,z_n) \mid (z_1,...,z_n)) - \max_{t\in[0,\infty]} U(\,t \mid z_n)\right| \leq \max_{t\in[0,\infty]}\left|U^N(t \mid (z_1,...,z_n)) - U(t \mid z_n)\right|.$$

By the uniform convergence of $U^N(t \mid (z_1,...,z_n))$ to $U(t \mid z_n)$, the right hand side converges to zero as $N\to\infty$, and therefore

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

Since $U(t \mid z_n)$ has a unique maximizer $t_n(z_n)$, we have

$$t_n^N(z_1,...,z_n) \to t_n(z_n).$$

■

Proof of Proposition 4. 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 Proposition 3 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,

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) \geq t_i(z_i). \tag{21}$$

Assume next that

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) > \lim_{N\to\infty} T_{i-1}^N\left(\frac{z_1}{N},...,\frac{z_{i-1}}{N}\right). \tag{22}$$

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, Proposition 3 implies that

$$\lim_{N\to\infty} T_i^N\left(\frac{z_1}{N},...,\frac{z_i}{N}\right) = t_i\left(z_i\right). \tag{23}$$

Combining equations (21) - (23), we have

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

For the player with the lowest signal, we have:

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

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

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

■

Proof of Theorem 2. We analyze the sequence of stopping times $\bar{t}_n\left(z_1,...,z_n\right)$, $n = 1, 2, ...,$ 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 strong law of large numbers, the sample average of n exponential random variables Z_1, $Z_2 - Z_1, ..., Z_n - Z_{n-1}$ converges almost surely to $1/g\left(0\left|\omega\right.\right)$ as $n \to \infty$. Assumption 4 implies that this identifies the true state ω. Therefore, the unconstrained stopping time $t_n\left(Z_n\right)$ defined in (8) converges to the first-best time as $n \to \infty$:

$$t_n\left(Z_n\right) \overset{a.s.}{\to} t^*\left(\omega\right). \tag{24}$$

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

$$\bar{t}_n\left(Z_1,...,Z_n\right) \overset{a.s.}{\to} \bar{t}_\infty, \tag{25}$$

and Equation (24) implies that

$$\lim_{n\to\infty} \Pr\left\{\bar{t}_n\left(Z_1,...,Z_n\right) \leq t\right\} = 0 \text{ for all } t < t^*\left(\omega\right). \tag{26}$$

Consider next the distribution of the first stopping time $t_1\left(Z_1\right)$. We have denoted the optimal stopping time under the lowest possible individual signal by $t\left(0\right)$. On the other hand, by assumption 4 we have $g\left(0\left|\omega\right.\right) > g\left(0\left|\max\Omega\right.\right)$ for any $\omega < \max\Omega$, and therefore the likelihood ratio across states ω and $\max\Omega$ goes to zero when $z_1 \to \infty$:

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

Therefore, we have

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

and hence:

$$\lim_{n \to \infty} \Pr\left\{\bar{t}_n(Z_1, ..., Z_n) < t(0)\right\} = 0, \tag{27}$$

and

$$\lim_{n \to \infty} \Pr\left\{\bar{t}_n(Z_1, ..., Z_n) > t\right\} > 0 \text{ for all } t < t^*(\max \Omega). \tag{28}$$

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 (25), 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. It then follows directly from (26), (27), and (28) that

$$F(t \mid \omega) = 0 \text{ for all } t < \max\{t(0), t^*(\omega)\} \text{ and}$$
$$F(t \mid \omega) < 1 \text{ for all } t < t^*(\max \Omega).$$

It remains to prove that

$$F(t \mid \omega) > 0 \text{ for all } t > \max\{t(0), t^*(\omega)\}.$$

Take an arbitrary $t > \max\{t(0), t^*(\omega)\}$. We want to show that $\Pr(t_n(z_n) < t \text{ for all } n) > 0$. Define for each $n = 1, 2...$ a cutoff value $\overline{z}_n$ as follows:

$$\overline{z}_n := \sup\{z : t_n(z) < t\},$$

so that $z_n < \overline{z}_n$ implies $t_n(z_n) < t$. Since $\lim_{z_1 \downarrow 0} t_1(z_1) = t(0) < t$, we have $\overline{z}_1 > 0$. Moreover, since $t_{n+1}(z) < t_n(z)$ for all z, $\{\overline{z}_n\}_{n=1}^{\infty}$ is a strictly increasing sequence. Since for any ω the sequence $\{t_n(Z_n)\}$ converges to $t^*(\omega)$ (almost surely), we have

$$\lim_{n \to \infty} (\overline{z}_n / n) = \overline{h} > 1/g(0 \mid \omega) := h^*.$$

Pick an arbitrary $h \in (h^*, \overline{h})$. We can then fix a large enough integer K such that the following implication holds:

$$\{z_n/n < h \text{ for } n > K\} \Longrightarrow z_n < \overline{z}_n. \tag{29}$$

Suppose that $z_n < \overline{z}_n$ for all $n = 1, ..., K$, and furthermore $z_K < Kh$. Since K is some fixed integer and $\{\overline{z}_n\}_{n=1}^{\infty}$ is a strictly increasing sequence with $\overline{z}_1 > 0$, this event has a strictly positive probability:

$$\Pr(t_n(z_n) < t \text{ for all } n = 1, ..., K, \text{ and } z_K < Kh) > 0. \tag{30}$$

Define the following process:

$$\{x_n\}_{n=0}^{\infty} := \{z_{K+n} - nh\}_{n=0}^{\infty}.$$

Note that $x_0 < 0$, and $\{x_n\}_{n=0}^{\infty}$ is a supermartingale:

$$\mathbb{E}(x_{n+1} | x_n) = z_{K+n} + h^* - (n+1)h < z_{K+n} - nh = x_n.$$

Moreover, since $z_K < Kh$, it follows from (29) that

$$\{x_n < 0 \text{ for } n \geq 1\} \Longrightarrow z_{K+n} < \overline{z}_{K+n}. \tag{31}$$

To show that $\Pr(t_n(z_n) < t \text{ for all } n > K) > 0$, we only need to show that

$$\Pr(x_n < 0 \text{ for all } n = 1, 2, ...) > 0.$$

To do that, we transform x_n to yield a bounded martingale g_n. Let $g_n = e^{\theta x_n}$, where θ is some constant. We have:

$$\begin{aligned}
\mathbb{E}(g_{n+1} | g_n) &= \mathbb{E}\left(e^{\theta(z_{K+n} + y_{K+n+1} - (n+1)h)} | x_n\right) = \mathbb{E}\left(e^{\theta(x_n + y_{K+n+1} - h)} | x_n\right) \\
&= e^{\theta x_n}\mathbb{E}\left(e^{\theta(y_{K+n+1} - h)} | x_n\right) = g_n \mathbb{E}\left(e^{\theta(y_{K+n+1} - h)} | x_n\right),
\end{aligned}$$

where $y_{K+n-1} \sim \exp(1/h^*)$. We want to choose θ so that g_n is a martingale:

$$\mathbb{E}\left(e^{\theta(y_{K+n+1} - h)} | x_n\right) = 1$$

or

$$\mathbb{E}\left(e^{\theta y_{K+n+1}} | x_n\right) = e^{\theta h}.$$

Noting that $y_{K+n-1} \sim \exp(1/h^*)$, recalling the moment generating function of exponential distribution, and taking log on both sides gives us:

$$\theta h = -\log(1 - \theta h^*).$$

Since $h > h^*$, this equation has a unique positive solution that we denote by $\theta^* > 0$. Hence

$$\{g_n\}_{n=0}^{\infty} := \left\{e^{\theta^* x_n}\right\}_{n=0}^{\infty}$$

is a martingale. Let n^* denote the stopping time:

$$n^* := \min\{n : x_n \geq 0\}.$$

For all $n < n^*$, we have $0 < g_n < 1$, so $\{g_n\}$ is a bounded martingale. Suppose that $\Pr(n^* < \infty) = 1$. Then by the Optional Sampling Theorem, we have

$$\mathbb{E}(g_{n^*} | x_0) = g_0 = e^{\theta x_0} < 1.$$

But if $n^* < \infty$, we have $x_{n^*} \geq 0$, so that $g_{n^*} > 1$. This is a contradiction, and we can conclude that

$$\Pr(n^* < \infty) = \Pr(x_n \geq 0 \text{ for some } n = 1, 2, ...) < 1,$$

or

$$\Pr(x_n < 0 \text{ for all } n = 1, 2, ...) > 0.$$

Combining this with (31) and (30) gives us

$$\Pr(t_n(z_n) < t \text{ for all } n = 1, 2, ...) > 0,$$

that is,

$$F(t \mid \omega) > 0\ .$$

■

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.

Billingsley, P. 1986. *Probability and Measure.* Wiley.

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.

Lang, S. 1983. *Real Analysis.* Addison-Wesley.

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, P. & J. Välimäki. 2011. "Learning and Information Aggregation in an Exit Game." *Review of Economic Studies* 78:1426–1462.

Murto, P. & J. Välimäki. 2012. "Large common value auctions with risk averse bidders.".

Pesendorfer, W. & J. Swinkels. 1997. "The Loser's Curse and Information Aggregation in Common Value Auctions." *Econometrica* 65(6):1247–1282.

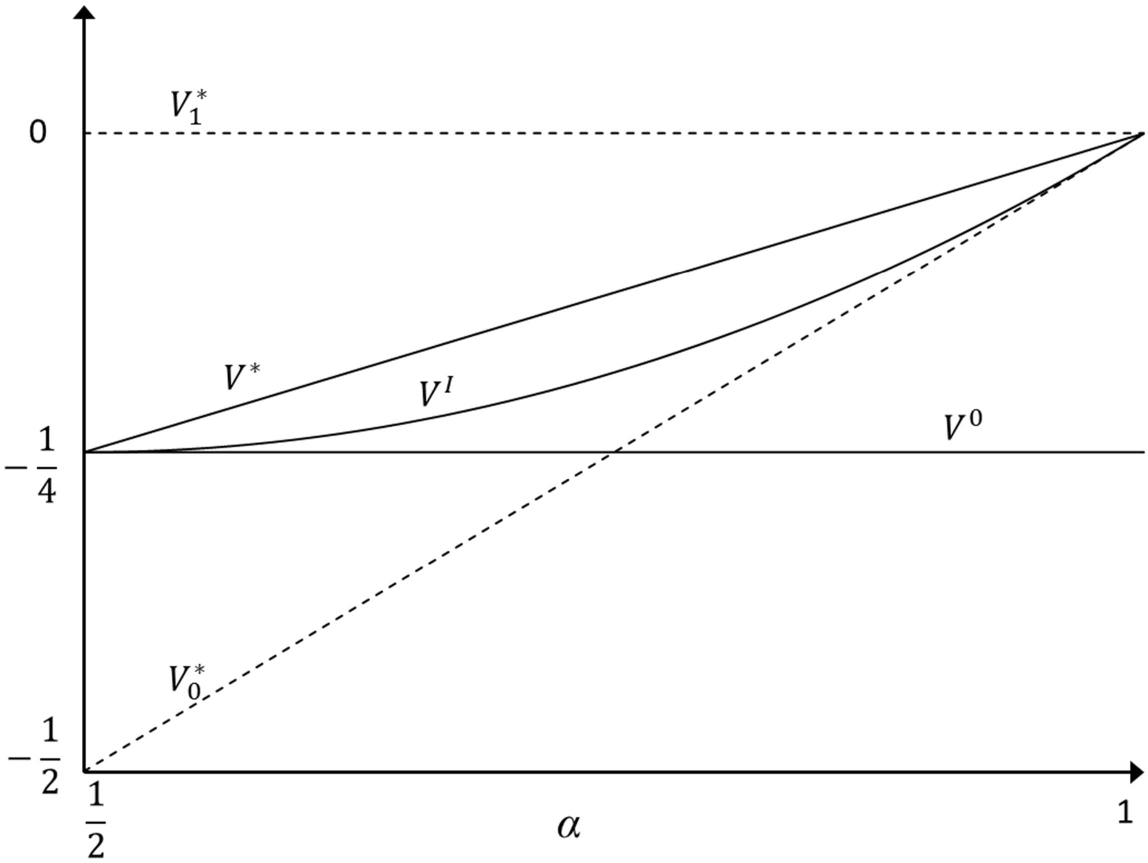

Figure 1: Payoffs as functions of signal precision in the quadratic binary example.

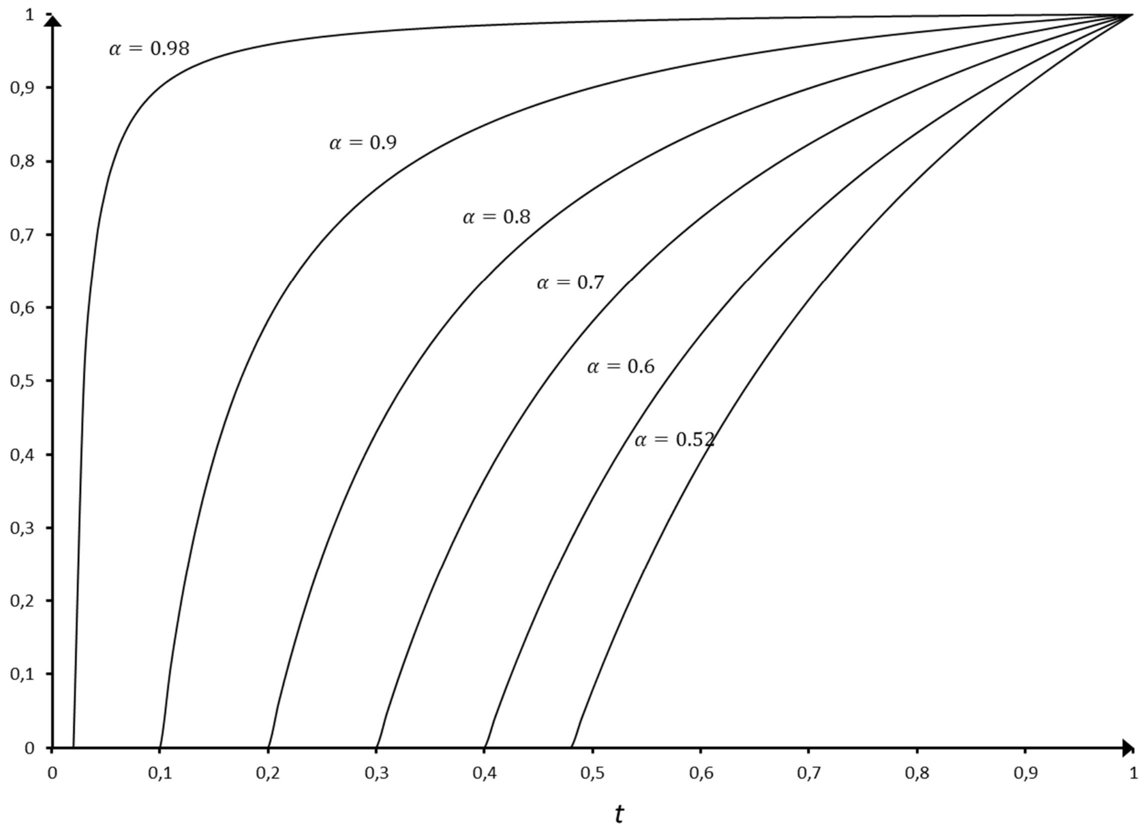

Figure 2: Probability distribution of the stopping time of the last player with various signal precisions ($\omega = 0$).

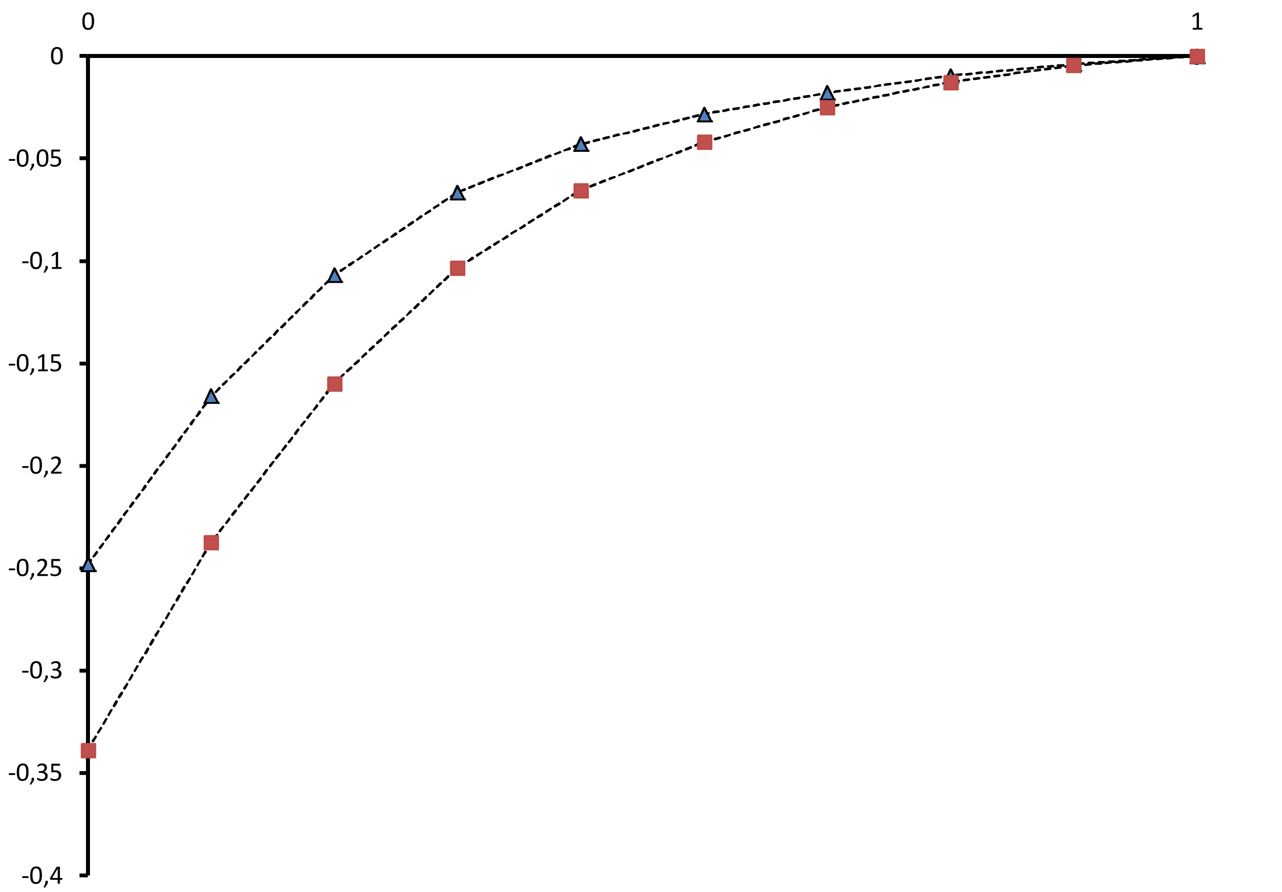

Figure 3: Equilibrium payoffs conditional on state in the ten-state example. Triangle marker: $\alpha = 1$. Square marker: $\alpha = 0.1$.

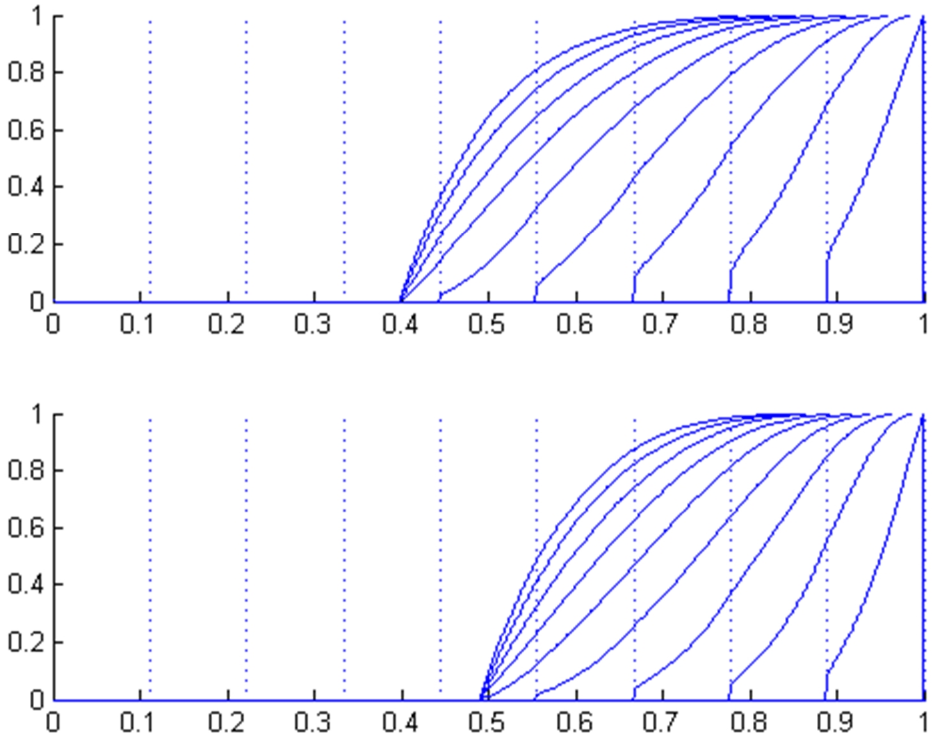

Figure 4: Conditional probability distributions of the stopping time of the last player in the ten-state example. Each solid curve corresponds to one state realization. Dashed lines correspond to first best stopping times for each state. Top panel: $\alpha = 1$. Bottom panel: $\alpha = 0.1$.