

Learning and Information Aggregation in an Exit Game^{*}

Pauli Murto[†] and Juuso Välimäki[‡]

First Version: December 2005

This Version: May 2008

Abstract

We analyze information aggregation in a stopping game with uncertain common payoffs. Players learn from their own private experiences as well as by observing the actions of other players. We give a full characterization of the symmetric mixed strategy equilibrium, and show that information aggregates in randomly occurring exit waves. Observational learning induces the players to stay in the game longer. The equilibria display aggregate randomness even for large numbers of players.

KEYWORDS: Learning, optimal stopping, dynamic games.

JEL CLASSIFICATION: C73,D81,D82,D83

^{*}We would like to thank numerous seminar audiences and, in particular, the editor Larry Samuelson, three anonymous referees, Dirk Bergemann, Hikmet Gunay, Godfrey Keller, Elan Pavlov and Peter Sorensen for useful comments. An earlier version of this paper was called “Learning in a Model of Exit”.

[†]Helsinki School of Economics, Department of Economics, P.O. Box 1210, 00101 Helsinki, Finland, pauli.murto@hse.fi

[‡]Helsinki School of Economics and University of Southampton, Department of Economics, P.O. Box 1210, 00101 Helsinki, Finland, juuso.valimaki@hse.fi

1 Introduction

Learning in dynamic decision problems comes in two different forms. Players learn from their own individual, and often private, observations about the fundamentals of their economic environment. At the same time, they may learn by observing the behavior of other players in analogous situations. In this paper, we analyze the interplay of these two modes of learning in an exit game with pure informational externalities.¹ We show that observational learning intensifies the effects of the fundamental uncertainty. All players stay longer in the game, and as a result those types that benefit from staying in the game win while those types that should exit lose. Even though private information accumulates steadily, it is revealed to the other players in occasional bursts.

There are a number of examples where both forms of learning are important. Learning about the quality of a service, the profitability of a new technology, or the size of a new market are examples of this type. In all these instances, it is reasonable to assume that part of the uncertainty is common to all agents and part is idiosyncratic. A new restaurant may be of high or low quality. A high quality restaurant is attractive to a larger fraction of the clientele than a low quality restaurant. Learning from others is useful to the extent that it can be used to determine whether the restaurant is good. It is not sufficient, however, if even a good restaurant does not appeal to all customers. Hence it is natural to consider a model where both private and observational learning are present.

To represent private learning, we use a standard discounted single-player experimentation model in discrete time. Each initially uninformed player collects information on her own binary type. The good types gain by staying in the game while the bad types gain by exiting the game. We assume that information accumulates according to a particularly simple form. Good types observe a perfectly informative signal with a constant probability in each period that they stay in the game while bad types never see any signals.² Uninformed players grow more pessimistic about their own type as time passes and the optimal strategy is simply to exit the game if uninformed after an optimally chosen number of periods.

Observational learning matters if a number of players face the same decision problem

¹The literature on strategic experimentation considers the third case where individual experiences are publicly observed. Examples of such models are Bolton & Harris (1999) and Keller, Rady & Cripps (2005). The focus in these papers is on the private provision of public information rather than on information aggregation.

²The actual form of information revelation is not very important for the logic of our model. The important assumption is that it takes time for even the most pessimistic individual player to exit the game.

and if their types are correlated. We model this correlation by assuming that there is a binary state of the world that determines the probability distribution of individual types. Conditional on the state, the players' types are identically and independently distributed. Whenever the exit decisions of a given player are sensitive to her information, her actions reveal information about her type and hence also about the state of the world. Uninformed players gain from additional information on the state, which creates an incentive to wait as in Chamley & Gale (1994). But in contrast to Chamley & Gale (1994), private learning makes it impossible for the players to delay indefinitely. Our model strikes a balance between the benefits from delaying in order to learn more from others and the costs from increased pessimism as a result of private learning.

We show that the game has a unique symmetric equilibrium in mixed strategies. The properties of this equilibrium become particularly sharp when we eliminate the effects of observation lags by reducing the time interval between consecutive decision moments. We show that the symmetric equilibrium can be characterized by two modes of behavior: In the *flow randomization mode*, bad news from no informative signals is balanced by the good news from the observation that no other player exits. Exits are infrequent and prior to any exit, the beliefs of the uninformed players evolve smoothly.

Following an exit, however, uninformed players become suddenly more pessimistic and have an incentive to exit the game. On the other hand, immediate exit by all uninformed players would release so much information that an individual player would find it optimal to wait. Hence the continuation play must be in mixed strategies where the randomizations are large enough to allow for a positive continuation payoff with positive probability. This leads immediately to further exits with a relatively high probability. We call this phase of consecutive exits an *exit wave*. An exit wave ends either in a collapse as the last uninformed player exits, or in a reversion to the flow randomization mode if there is a period with no exits. In the symmetric equilibrium, play fluctuates between these two modes until a collapse ends the game.³

When the number of players is increased towards infinity, the pooled information about aggregate state becomes accurate. A plausible conjecture would be that aggregate behavior conditional on state would become deterministic by the law of large numbers. We show that this is not the case. Even in the case with a large number of players, transitions between the phases remain random. The size of an individual exit wave as measured by the total number of exits during the wave also remains random. Information is thus

³Examples of models that display waves of action that resemble our exit waves include Bulow & Klemperer (1994) and Toxvaerd (2008). However, these models depend on the direct payoff externalities arising from scarcity.

aggregated during quick bursts, the exit waves, but the aggregation is not complete, since a random number of these bursts is needed to reveal all the information in the game. We derive explicit distributions for the exit probabilities during exit waves and hazard rates for their occurrence when the number of players is large.

The game has also asymmetric equilibria. In particular, there is an equilibrium in pure strategies, where players decide about their exit in a predetermined order conditioning their actions on the outcomes of the previous decisions. In this equilibrium those players that act later benefit from the information revealed by those who act earlier, and thus they have higher ex-ante expected payoffs.

Although asymmetric equilibria are perhaps not as relevant as the symmetric one due to coordination requirements (especially when the number of players is high), we nevertheless investigate what can be said about all equilibria of the game. Our main result confirms that all equilibria share a similar information aggregation property when the observation lag is small and the number of players is large. This result states that in the good state (where a higher fraction of players are successful), virtually all players exit at their optimal exit moment as if they knew the state in advance. In this sense, information is aggregated efficiently in the good state. But if the state is bad, information aggregation fails: players learn the true state too late, and as a result, they delay exit too much. The main message is that observational learning induces the players to stay in the game for longer than when acting alone. While this brings individual actions closer to ex-post efficient actions (good types are less likely to exit), it induces also a cost in terms of increased delay for the bad types.

This paper is related to the literature on herding and observational learning. That literature has studied the informational performance of games where players have private information at the beginning of the game. Much of this literature assumes an exogenously given order of moves for the players, e.g. Banerjee (1992), Bikhchandani, Hirshleifer & Welch (1992), and Smith & Sorensen (2000). This latter assumption has been relaxed by a number of papers. Among those, the most closely related to ours is Chamley & Gale (1994).⁴ In that paper a number of players consider investment in a waiting game that mirrors our setting. The model exhibits herding with positive probability: the players' beliefs may get trapped in an inaction region even when investment would be optimal. In our model the private learning during the game prevents the beliefs from getting trapped. The difference between the models is best seen by eliminating observation lags, i.e., letting period length go to zero. In Chamley and Gale, the length of the whole game goes to

⁴See also a more general model Chamley (2004). An early contribution along these lines is Mariotti (1992).

zero: information aggregates quickly but incompletely in one burst of investments that corresponds to one exit wave of our model. In our model, this limit involves smooth continuous time dynamics. All common values information eventually aggregates, but this takes time.

The two papers that are most closely related to ours in that they combine private and observational learning in a timing model are Caplin & Leahy (1994) and Rosenberg, Solan & Vieille (2007). While they are close to ours in their motivation, each makes a crucial modeling assumption that leads to qualitatively different information aggregation properties to ours. Caplin and Leahy assume a continuum of players from the beginning. This assumption rules out what is a key property of our model: a large number of players can jointly release a small amount of information. Rosenberg, Solan & Vieille (2007) work with a finite number of players like we do, but they assume signals that make some players so pessimistic after one period that exiting is the dominant strategy. As a result, when the number of players is increased, the exit behavior after the first period reveals the state by the law of large numbers. Due to these modeling assumptions, the aggregate behavior in the large game limit is essentially deterministic conditional on state both in Caplin & Leahy (1994) and Rosenberg, Solan & Vieille (2007). Our model complements these papers by showing that information may also aggregate slowly through randomly occurring exit waves, even when the pooled information is precise.

Finally, by combining common and idiosyncratic uncertainty, our paper relaxes the assumption of perfect payoff correlation across players made in Chamley & Gale (1994), Caplin & Leahy (1994), and Rosenberg, Solan & Vieille (2007). The pure common values case is obtained in our model as a limiting case.

The paper is organized as follows. Section 2 sets up the discrete time model. Section 3 provides the analysis of the symmetric and pure strategy equilibria of the model. In Section 4, we characterize the symmetric equilibrium explicitly in the continuous time limit. In Section 5 we prove that all equilibria aggregate information approximately efficiently in the good state when the number of players is large. Section 6 concludes.

2 Model

The model is in discrete time with periods $t = 0, 1, \dots, \infty$. The discount factor per period is $\delta = e^{-r\Delta}$, where Δ is the length of a period. The set of players is denoted by $\mathcal{N} = \{1, \dots, N\}$.

Before the game starts, nature chooses the (aggregate) state randomly from two alternatives: $\theta = \theta_H$ (high) and $\theta = \theta_L$ (low). Let p^0 denote the common prior $p^0 =$

$\Pr(\theta = \theta_H)$. After choosing the state, nature chooses randomly and independently the individual type for each player. Each player is either good or bad. If $\theta = \theta_H$, the probability of being good is α , while if $\theta = \theta_L$, the probability of being good is β , where $0 < \beta < \alpha < 1$. The player types are drawn independently for all players. All types are initially unobservable to all players, but the parameters p^0 , α , and β are common knowledge.

The information about nature's choices arrives gradually during the game as follows. In each period, each player gets a random signal $\zeta \in \{0, 1\}$. Signals have two functions: they transmit information and payoffs. For a player of bad type, $\zeta = 0$ with probability 1. For a good player, $\Pr(\zeta = 1) = \lambda\Delta$, where λ is a commonly known parameter. The signal realizations across periods and players (conditional on the state and the type) are independent. We call the signal $\zeta = 1$ a *positive* signal, since it entails a positive payoff (see next paragraph) and reveals to a player that her type is good. Each player observes only her own signals. We use the terms *informed* and *uninformed* to refer to the players' knowledge of their own type: players who have had a positive signal are informed, other players are uninformed.

At the beginning of each period, all active players make a binary decision: stop or continue. Stopping is costless, but irreversible: once a player stops, she becomes inactive and receives the outside option payoff 0. If the player continues, she pays the (opportunity) cost $c \cdot \Delta$, observes a signal $\zeta \in \{0, 1\}$ that yields payoff $\zeta \cdot v$, and then moves to the next period. Here c and v are parameters for which $c < \lambda v$. Since we assume risk neutrality, the payoff per period is $(\lambda v - c) \Delta > 0$ for a good player and $-c\Delta < 0$ for a bad player. As a consequence, bad types want to stop and good types never stop.

Within each period the players act simultaneously, but they know each other's previous actions. However, they do not observe each others' signals, and therefore they do not know whether the others are informed or uninformed. Note that new information arrives to the players via two channels: their own signals and observations on other players' behavior. In the terminology of learning models, they engage simultaneously in experimentation and observational learning.

The history of player i consists of the private history recording her own signal history, and the public history recording the actions of all the players. Since observing a positive signal reveals fully the player's type, the only thing that matters in each private history is whether it contains at least one positive signal. Since it is always a strictly dominant action for any informed player to continue, we can take it as given that informed players never stop. Strategies are therefore fully described by the stopping behavior of the uninformed players. For those, the only relevant history is the public history, and from

now on we call this simply the *history*. Denote the history in period t by h^t and define it recursively as follows:

$$\begin{aligned} h^0 &= \emptyset, \\ h^t &= h^{t-1} \cup a^{t-1} \quad \forall t \in \{1, 2, \dots\}, \end{aligned}$$

where $a^t = (a_1^t, \dots, a_N^t)$ is a vector where each $a_i^t \in \{0, 1\}$ denotes an indicator for i continuing at period t .

Denote by H^t the set of all possible histories up to t and let $H = \bigcup_{t=0}^{\infty} H^t$. Since stopping is irreversible, $a_i^t = 0$ implies that $a_i^{t'} = 0$ for all $t' > t$ for all elements of H . Denote by $H_i \equiv \{h^t \in H \mid a_i^{t-1} = 1\}$ the set of histories, in which i has not yet stopped. Denote by $A(h^t) \equiv \{i \in \mathcal{N} \mid h^t \in H_i\}$ the set of active players, and let $n(h^t)$ denote their number.

A strategy for an uninformed player i is a mapping

$$\sigma_i : H_i \rightarrow [0, 1]$$

that maps all histories where i is active to a stopping probability. The strategy profile is $\sigma = (\sigma_1, \dots, \sigma_N)$.

The value of a player is the expected discounted sum of future cash flows as estimated on the basis of her own signal history, observations of other players' behavior, and initial prior probability p^0 . Denote by $V_i(h^t; \sigma)$ the value of an uninformed player i after history h^t and with profile σ . The value of an informed player is easy to calculate explicitly:

$$V^+ = \frac{(\lambda v - c) \Delta}{1 - \delta}.$$

By equilibrium, we mean a Perfect Bayesian Equilibrium of the above game. In an equilibrium, all actions in the support of $\sigma_i(h^t)$ are best responses to σ_{-i} for all i and for all h^t .

2.1 Beliefs

Players have two types of posterior beliefs. First, they have a belief concerning the state. This is directly relevant, because it determines the posterior on their own type. Second, they have beliefs concerning the information of other players. This is indirectly relevant, because it determines what can be inferred about the state by observing others' behavior.

Given h^t and σ , all beliefs can be summarized in three quantities: $p_i(h^t, \sigma)$, $q_H^j(h^t, \sigma)$, and $q_L^j(h^t, \sigma)$, $i, j \in \mathcal{N}$. Here $p_i(h^t, \sigma)$ is the belief of player i about the state (i.e.

probability with which $\theta = \theta_H$), and $q_\theta^j(h^t, \sigma)$, $\theta \in \{\theta_H, \theta_L\}$, is the posterior held by any player (or equivalently, an outside observer) on the event " j is informed", conditional on state. Given these, we may write the unconditional posterior held by i about the event " j is informed" as:

$$q_i^j(h^t, \sigma) = q_H^j(h^t, \sigma) p_i(h^t, \sigma) + q_L^j(h^t, \sigma) (1 - p_i(h^t, \sigma)). \quad (1)$$

An individual player is ultimately concerned about her own type rather than the state. Denote by $s_i(h^t, \sigma)$ the belief of i on her own type, i.e.

$$s_i(h^t, \sigma) = \Pr\{\text{type of } i \text{ is "good"} \mid h^t, \sigma\}.$$

Note that conditional on state, the belief of every uninformed player of her own type is identical (this is because all uninformed have an identical private history, which is all that matters conditional on state). Let s_θ^t denote the belief of any uninformed player of her own type, conditional on θ . Applying Bayes' rule to a private history containing no positive signals gives:

$$s_\theta^t = \frac{s_\theta^0 (1 - \lambda \Delta)^t}{s_\theta^0 (1 - \lambda \Delta)^t + 1 - s_\theta^0}, \quad (2)$$

where $\theta \in \{\theta_H, \theta_L\}$, and $s_H^0 = \alpha$, $s_L^0 = \beta$. Using this, we may now express i 's belief of her own type in terms of her belief of state:

$$\begin{aligned} s_i(h^t, \sigma) &= p_i(h^t, \sigma) s_H^t + (1 - p_i(h^t, \sigma)) s_L^t \\ &= p_i(h^t, \sigma) (s_H^t - s_L^t) + s_L^t. \end{aligned} \quad (3)$$

We may also express a player's beliefs of state and other players' information conditional on her own type. Denote by $p_{i+}(h^t, \sigma)$ and $p_{i-}(h^t, \sigma)$ player i 's posterior probability on the good state, conditional on being herself of good (+) or bad (-) type, respectively:

$$p_{i+}(h^t, \sigma) = \frac{s_H^t p_i(h^t, \sigma)}{s_H^t p_i(h^t, \sigma) + s_L^t (1 - p_i(h^t, \sigma))}, \quad (4)$$

$$p_{i-}(h^t, \sigma) = \frac{(1 - s_H^t) p_i(h^t, \sigma)}{(1 - s_H^t) p_i(h^t, \sigma) + (1 - s_L^t) (1 - p_i(h^t, \sigma))}. \quad (5)$$

Using these, i 's belief of j 's information conditional on her own type being good and bad, respectively, can be expressed as:

$$q_{i+}^j(h^t, \sigma) = q_H^j(h^t, \sigma) p_{i+}(h^t, \sigma) + q_L^j(h^t, \sigma) (1 - p_{i+}(h^t, \sigma)), \quad (6)$$

$$q_{i-}^j(h^t, \sigma) = q_H^j(h^t, \sigma) p_{i-}(h^t, \sigma) + q_L^j(h^t, \sigma) (1 - p_{i-}(h^t, \sigma)). \quad (7)$$

2.2 Learning

Within each period, the beliefs react to two random events. First, at the beginning of the period, other players' actions give rise to observational learning. Second, during the period, private signals induce another belief update.

We first derive the update in the belief about state. Let $\tilde{p}_i(h^t, \sigma) \equiv \log\left(\frac{p_i(h^t, \sigma)}{1-p_i(h^t, \sigma)}\right)$ denote the log-likelihood ratio of i 's belief of state. At the beginning of the period players observe each others' exit decisions. Player j exits with probability $\sigma_j(h^t)$. This induces the following change in i 's belief:

$$\Delta_j \tilde{p}_i(a_j) = \begin{cases} \log\left(\frac{1-(1-q_H^j)\sigma_j(h^t)}{1-(1-q_L^j)\sigma_j(h^t)}\right), & \text{if } a_j^t = 1 \\ \log\left(\frac{(1-q_H^j)\sigma_j(h^t)}{(1-q_L^j)\sigma_j(h^t)}\right), & \text{if } a_j^t = 0 \end{cases}. \quad (8)$$

The total change in i 's belief from observing the actions of all $\mathcal{N} \setminus i$ other players is then:

$$\Delta^1 \tilde{p}_i(a^t) = \sum_{j \in \mathcal{N} \setminus i} \Delta_j \tilde{p}_i(a_j^t). \quad (9)$$

After this observational learning, players obtain private signals during the period. The change in \tilde{p} from this is identical for all uninformed players, who get a non-positive signal and thus remain uninformed:

$$\Delta^2 \tilde{p} = \log\left(\frac{1-s_H^t \lambda \Delta}{1-s_L^t \lambda \Delta}\right). \quad (10)$$

Combining the two sources of learning, we have:

$$\tilde{p}_i(h^{t+1}, \sigma) = \tilde{p}_i(h^t, \sigma) + \Delta^1 \tilde{p}_i(a^t) + \Delta^2 \tilde{p}, \quad (11)$$

and the new belief for a player that remains uninformed can be written as:

$$p_i(h^{t+1}, \sigma) = \frac{\exp(\tilde{p}_i(h^{t+1}, \sigma))}{1 + \exp(\tilde{p}_i(h^{t+1}, \sigma))}. \quad (12)$$

Consider next the updates in a player's beliefs over another player's information, $q_\theta(h^t, \sigma)$. First, at the beginning of the period, players observe each other's stopping behavior. If player j continues, then other players' belief of j 's information changes to $\tilde{q}_\theta^j(h^t, \sigma)$:

$$\tilde{q}_\theta^j(h^t, \sigma) = \frac{q_\theta^j(h^t, \sigma)}{1 - \sigma_j(h^t)(1 - q_\theta^j(h^t, \sigma))}, \quad \theta \in \{\theta_H, \theta_L\}. \quad (13)$$

Second, players understand that other players may become informed within the current period (after exit decisions have been undertaken), which induces an additional update. After this second update, the new belief is:

$$\begin{aligned}
q_\theta^j(h^{t+1}, \sigma) &= \widehat{q}_\theta^j(h^t, \sigma) + (1 - \widehat{q}_\theta^j(h^t, \sigma)) s_\theta^t \lambda \Delta \\
&= s_\theta^t \lambda \Delta + \frac{q_\theta^j(h^t, \sigma) (1 - s_\theta^t \lambda \Delta)}{1 - \sigma_j(h^t) (1 - q_\theta^j(h^t, \sigma))}, \theta \in \{\theta_H, \theta_L\}.
\end{aligned} \tag{14}$$

3 Equilibrium Analysis

3.1 Isolated player

As a useful starting point, we consider an isolated player that can only learn from her own signals. This player faces a standard stopping problem. Denote by s the current belief of an uninformed player about her type. If the player continues for another period, but still receives no positive signal, the new posterior $s + \Delta s$ is obtained by Bayes' rule:

$$s + \Delta s = \frac{s(1 - \lambda \Delta)}{s(1 - \lambda \Delta) + 1 - s} = \frac{1 - \lambda \Delta}{s^{-1} - \lambda \Delta}. \tag{15}$$

Denote the value function of an isolated player by $V_m(s)$. If it is optimal to stop, this value is 0. If the player continues, she gets a positive signal with probability $s\lambda\Delta$ in which case the value jumps to $V_m(1) = V^+$. Without a positive signal s falls to $s + \Delta s$. Bellman's equation can thus be written as:

$$V_m(s) = \max \left\{ 0, \widetilde{V}_m(s) \right\}, \tag{16}$$

where $\widetilde{V}_m(s)$ is the value of continuing for at least one more period:

$$\widetilde{V}_m(s) \equiv -c\Delta + s\lambda\Delta (v + \delta V^+) + (1 - s\lambda\Delta) \delta V_m \left(\frac{1 - \lambda\Delta}{s^{-1} - \lambda\Delta} \right). \tag{17}$$

The optimal policy is to stop as soon as s falls below some threshold level, denoted s^* . The value function $V_m(s)$ must be weakly increasing in s . By (17), $\widetilde{V}_m(s)$ is then strictly increasing in s . The threshold s^* is obtained from (17) by setting $\widetilde{V}_m(s^*) = 0$ and noting that $V_m(s) = 0$ for $s < s^*$:

$$s^* = \frac{c}{\lambda(v + \delta V^+)}. \tag{18}$$

We shall see that s^* plays a crucial role also in the model with many players. Denote by t^* the period at which s falls below s^* in case there is no positive signal.

The analysis of the isolated player leads to our first result concerning N players. This result, valid for all equilibria, says that after any history, any player is at least as well off as an isolated player would be (given the same current belief of her type), but it is not possible that *all* players are strictly better off:

Lemma 1 *Let σ be an equilibrium profile. For any $h^t \in H$, $V_i(h^t; \sigma) \geq V_m(s_i(h^t; \sigma))$ for all $i \in A(h^t)$ and $V_i(h^t; \sigma) = V_m(s_i(h^t; \sigma))$ for some $i \in A(h^t)$. Further, $\sigma_i(h^t) = 0$ whenever $s_i(h^t, \sigma) > s^*$.*

Since $s_i(h^t, \sigma) > s^*$ for all $t < t^*$, all players stay with probability one until time t^* . Since the players reveal information only through their exit decisions, there can never be any information sharing before time t^* .

3.2 Symmetric equilibrium

A profile σ is symmetric if $\sigma_i(h^t) = \sigma_j(h^t) \equiv \sigma(h^t)$ for all $i, j \in A(h^t)$ and for all h^t . When σ is symmetric, all uninformed players update their beliefs in the same way, and therefore they all share common posteriors. We thus omit all sub- and superscripts referring to individual players throughout this section.

With a symmetric σ , after history h^t , a given player observes the actions of $n(h^t) - 1$ other players that stop with probability $\pi = \sigma(h^t)$. It is simplest to work here directly with the beliefs concerning the players' own type. In this case, the inference from other players' actions can be defined in terms of $q_+ = q_+(h^t, \sigma)$ and $q_- = q_-(h^t, \sigma)$, given in (6) and (7).⁵ The number of other players that exit follows a binomial distribution, where the exit probability for each individual player is $(1 - q_+) \pi$ if the player's own type is good, and $(1 - q_-) \pi$ otherwise. Denoting by s the current belief of an uninformed player of her own type, the new belief directly after observing k out of $n - 1$ other players stopping is:

$$s'(k; n, \pi, s, q_+, q_-) = \frac{s [(1 - q_+) \pi]^k [1 - (1 - q_+) \pi]^{n-1-k}}{s [(1 - q_+) \pi]^k [1 - (1 - q_+) \pi]^{n-1-k} + (1 - s) [(1 - q_-) \pi]^k [1 - (1 - q_-) \pi]^{n-1-k}}. \quad (19)$$

In any symmetric equilibrium, all uninformed players have the same expected payoff. It follows from Lemma 1 that this payoff must equal the value of an isolated player. Thus, to make a player indifferent between exiting and staying, the observational learning in the current period must suffice to make continuation value zero even when ignoring any observational learning that might take place in the future. To formalize this key property of the symmetric equilibrium, we define an auxiliary value function that we call the *single-observation continuation value*. This is the continuation payoff of a player that observes the randomization of other players in the current period, but after this period will behave as an isolated player:

⁵The beliefs $q_+(h^t, \sigma)$ and $q_-(h^t, \sigma)$ are updated using (12), (14), and (4) - (7).

$$\begin{aligned}
C_n^\pi(s, q_+, q_-) &= \mathbb{E} \tilde{V}_m(s') \\
&= \sum_{k=0}^{n-1} \left[\binom{n-1}{k} [(1-q)\pi]^k [1 - (1-q)\pi]^{n-1-k} \tilde{V}_m(s') \right], \quad (20)
\end{aligned}$$

where $q = sq_+ + (1-s)q_-$ and $s' = s'(k; n, \pi, s, q_+, q_-)$ as given by (19).

The following proposition states that there is a unique symmetric equilibrium in mixed strategies. Whenever s is above s^* , it is a dominant action for them to continue. When s is below a certain lower threshold $\underline{s}(n^t, q_+^t, q_-^t)$, then it is a dominant action to exit. Between these thresholds, there is a unique stopping probability $\pi^*(n^t, s^t, q_+^t, q_-^t)$ that releases just enough information to make uninformed players indifferent between stopping and continuing:

Theorem 1 *There is a unique symmetric equilibrium $\sigma^S = [\sigma_1^S, \dots, \sigma_N^S]$ defined by:*

$$\sigma_i^S(h^t) = \begin{cases} 0 & , \text{ if } s^t \geq s^* \\ \pi^*(n^t, s^t, q_+^t, q_-^t) & , \text{ if } \underline{s}(n^t, q_+^t, q_-^t) \leq s^t < s^* \\ 1 & , \text{ if } s^t < \underline{s}(n^t, q_+^t, q_-^t) \end{cases} \quad , i \in A(h^t), \quad (21)$$

where $n^t = n(h^t)$, $s^t = s(h^t, \sigma^S)$, $q_+^t = q_+(h^t, \sigma^S)$, $q_-^t = q_-(h^t, \sigma^S)$, and where $\underline{s}(\cdot)$ and $\pi^*(\cdot)$ are the unique values implicitly defined by equations:

$$\underline{s}(n, q_+, q_-) = \{s \in (0, s^*]; C_n^1(s, q_+, q_-) = 0\}, \quad (22)$$

$$\pi^*(n, s, q_+, q_-) = \{\pi \in (0, 1]; C_n^\pi(s, q_+, q_-) = 0\}. \quad (23)$$

Note that if $s^t < s^*$, then $\tilde{V}_m(s^t) < 0$. To have $C_n^\pi(s^t, q_+^t, q_-^t) \geq 0$, there must then be a positive probability that $s^{t+1} > s^*$. Since $s(h^{t+1}; \sigma^S)$ is at its highest after histories h^t where no player exits, the following must hold in equilibrium:

Remark 1 *If $\underline{s}(n(h^t), q_+(h^t; \sigma^S), q_-(h^t; \sigma^S)) \leq s(h^t; \sigma^S) < s^*$ and no player exits at h^t , then $s(h^{t+1}; \sigma^S) > s^*$.*

The symmetric equilibrium path can be verbally described as follows. Given that $s^0 > s^*$, all uninformed players initially stay with probability one until s^t falls below s^* . At that point they start to randomize. In each period, the remaining uninformed players update their current beliefs after observing the number of exits. Positive news may move s^t back above s^* , in which case randomizations end until s^t falls again below s^* . When s^t falls below $\underline{s}(n^t, q_+^t, q_-^t)$, the remaining uninformed players become so pessimistic that

the information held by others, even if it were fully shared, is not sufficiently valuable to counterweight the cost of staying an extra period. At that point all remaining uninformed players exit. A definite bound for when the game must be over is when s_H^t falls below s^* : then even knowing for sure that $\theta = \theta_H$ would not justify continuing another period (we have then $\underline{s}(n^t, q_+^t, q_-^t) = s^*$). In Section 4 we give a full description of the symmetric equilibrium path in the limit as $\Delta \downarrow 0$.

3.3 Asymmetric equilibrium in pure strategies

The model has asymmetric equilibria in addition to the symmetric equilibrium discussed above. We show next that there is a class of pure strategy equilibria that gives a higher ex-ante expected sum of payoffs to the players. A strategy profile σ is a pure strategy profile if $\sigma_i(h^t) \in \{0, 1\}$ for all $h^t \in H$, $i \in A(h^t)$. When a pure strategy commands i to stop, other players learn perfectly i 's information. If i exits, she becomes inactive; if she stays, her type becomes common knowledge. Given a pure strategy σ , let $\tilde{A}(h^t, \sigma) \subseteq A(h^t)$ denote the set of players, who have not yet revealed their information. We call $\tilde{A}(h^t, \sigma)$ the set of *informative* players and $\tilde{n}(h^t, \sigma) \leq n(h^t)$ is the number of informative players.

Given s , q_+ , and q_- , let $\varphi(s, q_+, q_-)$ denote the minimum number of informative players required to reveal their private history in order to make continuation without further observational learning optimal for an uninformed player:

$$\varphi(s, q_+, q_-) \equiv \begin{cases} 0 & \text{if } \tilde{V}_m(s) > 0, \\ \min(n \in \{1, 2, \dots\} \mid C_{n+1}^1(s, q_+, q_-) \geq 0) & \text{if } \tilde{V}_m(s) \leq 0 \end{cases}, \quad (24)$$

where $C_n^1(\cdot)$ is the single-observation continuation value with stopping probability one. If $\lim_{n \rightarrow \infty} C_n^1(s, q_+, q_-) \leq 0$, then we define $\varphi(s, q_+, q_-) = \infty$.

The next Theorem says that there exists a set of pure strategy equilibria defined by φ :

Theorem 2 *Let σ^P be a pure-strategy profile for which*

$$\#\left\{i \in \tilde{A}(h^t, \sigma^P) \mid \sigma_i^P(h^t) = 1\right\} = \min[\varphi(s^t, q_+^t, q_-^t), \tilde{n}(h^t, \sigma^P)] \quad (25)$$

for each $h^t \in H$, where $s^t = s(h^t, \sigma^P)$, $q_+^t = q_+(h^t, \sigma^P)$, $q_-^t = q_-(h^t, \sigma^P)$. Then σ^P is an equilibrium.

The pure strategy equilibrium assigns to each history a set (possibly empty) of players that reveal their private histories. Equations (24) and (25) define the number of those

players so that their information makes the single-observation continuation value positive for those who stay, but not for those who exit. The proof shows that if the single-observation continuation value is negative for a player, then a deviation by mimicking an informed player cannot make continuation profitable. Even though such a deviation gives access to information that others reveal later, the deviation makes other players overly optimistic and future information flow is delayed.

While there are many strategy profiles satisfying Theorem 2, they are all obtained via a permutation of the players' identities. A notable difference between a pure strategy equilibrium and the symmetric equilibrium is that the former gives a higher ex-ante total payoff. The expected payoff in the symmetric equilibrium coincides with the stand-alone payoff. In contrast, the pure strategy equilibrium payoff of those players that exit after observing other players' decisions is higher.

We now turn to the characterization of the symmetric equilibrium in the continuous time limit. In Section 5 we will discuss the properties of all equilibria.

4 Exit Waves

In this section, we characterize the symmetric equilibrium in the limit as $\Delta \downarrow 0$. We do this for two reasons. First, we want to rule out any effects that observation lags might have on equilibrium properties. Second, this limit reveals sharply the inherent dynamics of the model: information aggregation happens in randomly occurring exit waves.

It is convenient to express the limiting properties of the equilibrium in continuous time. This is indicated by using argument t in parenthesis. Let us stress, however, that our model is not defined in continuous time: strategies, beliefs, etc. are only defined for histories consisting of a finite number of periods. Therefore, the continuous time equations are to be understood merely as a representation of the belief dynamics of the unique symmetric equilibrium as $\Delta \downarrow 0$.

4.1 Equilibrium Path Prior to Exit

Let $[t, t + dt]$ denote an arbitrary fixed time interval of real time. We let first $\Delta \downarrow 0$, so that the number of decision moments within $[t, t + dt]$ goes to infinity. In the second step, we let $dt \downarrow 0$ to arrive at a continuous time description of the equilibrium path.

We start with some limiting properties of the symmetric equilibrium. In expression (18), s^* , depends on Δ through δ and V^+ . To emphasize this dependence, we write it here as $s^*(\Delta)$. With a slight abuse of notation, we use s^* without an argument in this

section to refer to the exit threshold in the limit $\Delta \downarrow 0$.

Lemma 2 Denote by $\pi^*(n, s, q_+, q_-; \Delta)$ the equilibrium exit probability defined in Theorem 1, given period length Δ . Let $n \geq 2$ and $q_- < q_+ < 1$. Then, we have

$$\lim_{\Delta \downarrow 0} \pi^*(n, s^*(\Delta) - f(\Delta), q_+, q_-; \Delta) = 0,$$

where $f(\Delta)$ is an arbitrary positive function for which $\lim_{\Delta \downarrow 0} f(\Delta) = 0$.

We consider next the time path of s in the symmetric equilibrium. Whenever s crosses $s^*(\Delta)$, its value in the next period must be $s > \frac{1-\lambda\Delta}{(s^*(\Delta))^{-1}-\lambda\Delta}$. Lemma 2 implies that the equilibrium exit probability after any history where s has just crossed $s^*(\Delta)$ goes to zero as $\Delta \downarrow 0$. On the other hand, Theorem 1 implies that when no player exits, the next period belief s' must be above $s^*(\Delta)$. As soon as s has fallen below $s = s^*$, the players randomize with a probability that bounces s immediately back above $s^*(\Delta)$ (if no player exits). After this, s drifts down until it crosses $s^*(\Delta)$ again at which point another randomization bounces it back above $s^*(\Delta)$, and so on. As Δ gets smaller, the band around $s^*(\Delta)$ within which s drifts narrows down, and in the limit $\Delta \downarrow 0$, s must stay infinitely close to s^* .

Next we fix the interval $[t, t + dt]$, and consider the belief of an arbitrary uninformed player, who observes neither a positive signal nor an exit by another player within the interval. Since her belief remains (arbitrarily) close to s^* throughout the interval, the belief updates from seeing no positive signals and no exits must offset each other. Since the probability of a positive signal is of the order $O(dt)$, the same must be true for the probability of observing an exit. Therefore, when $dt \downarrow 0$, we may denote by $\rho(t) \cdot dt$ the probability with which an arbitrary uninformed player exits within $[t, t + dt]$, where $\rho(t)$ is to be interpreted as the hazard rate of exit of an arbitrary uninformed player. We call $\rho(t)$ the equilibrium flow rate of exit.

Adapting equations (8) - (12) appropriately and calculating $\lim_{dt \downarrow 0} \frac{p(t+dt) - p(t)}{dt}$ leads to:⁶

$$\dot{p}(t) = [-\lambda(s_H(t) - s_L(t)) + (n-1)\rho(t)(q_H(t) - q_L(t))]p(t)(1-p(t)). \quad (26)$$

Similarly, adapting (13) - (14) to calculate the limit $\lim_{dt \downarrow 0} \frac{q_\theta(t+dt) - q_\theta(t)}{dt}$, we get:

$$\dot{q}_\theta(t) = [\lambda s_\theta(t) + \rho(t)q_\theta(t)][1 - q_\theta(t)]. \quad (27)$$

⁶Equations (8) - (12) give belief updates within a single period. Here we want to derive the corresponding updates for an interval $[t, t + dt]$ that contains an infinite number of periods. However, when dt is small, the only change required is to replace Δ by dt , and $\sigma_j(t)$ by $\rho(t) \cdot dt$.

Letting $\Delta \downarrow 0$ in (2) gives:

$$\dot{s}_\theta(t) = -\lambda s_\theta(t) (1 - s_\theta(t)). \quad (28)$$

Differentiating (3), and using (26) and (28), the dynamics of $s(t)$ can be written as:

$$\begin{aligned} \dot{s}(t) &= \dot{p}(t) (s_H(t) - s_L(t)) + p(t) (\dot{s}_H(t) - \dot{s}_L(t)) + \dot{s}_L(t) \\ &= -\lambda s(t) [1 - s(t)] + [n - 1] \rho(t) p(t) [1 - p(t)] [s_H(t) - s_L(t)] [q_H(t) - q_L(t)] \end{aligned} \quad (29)$$

Equations (26) - (29) give the belief dynamics with an arbitrary $\rho(t)$. The equilibrium requirement is that $\rho(t)$ keeps $s(t)$ at constant value s^* as long as no player exits. Requiring $\dot{s}(t) = 0$ in (29) gives:

$$\rho(t) = \frac{\lambda s(t) (1 - s(t))}{(n - 1) (q_H(t) - q_L(t)) (s_H(t) - s_L(t)) p(t) (1 - p(t))} \equiv \frac{\Gamma(t)}{n - 1}. \quad (30)$$

Note that $s(t) = s^*$ is equivalent to $p(t) = p^*(t)$, where

$$p^*(t) \equiv \frac{s^* - s_L(t)}{s_H(t) - s_L(t)}.$$

Note, however, that $s(h^t, \sigma)$ is bounded from above by s_H^t for any h^t and σ . Therefore when s_H^t falls below s^* , it becomes impossible to keep an uninformed player indifferent. To understand what happens at this point, let

$$\bar{p}(p, q_H, q_L, n) \equiv \frac{q_H^{n-1} \cdot p}{q_H^{n-1} \cdot p + q_L^{n-1} \cdot (1 - p)}$$

denote the maximum belief on the state of the world that information sharing could induce. If all the other players turn out to be informed, then $\bar{p}(p, q_H, q_L, n)$ is the belief of the n :th player, who is uninformed. Assuming that the players continue randomizing according to (30) and still assuming no exits, $p(t)$, $q_H(t)$, and $q_L(t)$ keep increasing, and eventually $p(t)$ and $\bar{p}(p, q_H, q_L, n)$ cross each other. Close to that moment, $p(t)$ approaches $\bar{p}(p, q_H, q_L, n)$, and $\pi(t)$ goes to infinity. At the same time $q(t)$, $q_H(t)$, and $q_L(t)$ approach one. This means that by the time when p^* and $\bar{p}(p, q_H, q_L, n)$ cross, all uninformed players have already exited with probability 1, and at that moment it becomes common knowledge that all the remaining players are informed and the game is over.

If the number of players goes to infinity, their pooled information is sufficient to reveal the true state of nature to an arbitrary degree of precision. This means that $\bar{p}(p, q_H, q_L, n) \xrightarrow{n \rightarrow \infty} 1$ whenever $p > 0$, $q_L < q_H$. Thus, as $N \rightarrow \infty$, the game can continue as long as

$$p^*(t) \equiv \frac{s^* - s_L(t)}{s_H(t) - s_L(t)} < 1.$$

The latest moment when the game must end is when $s_H(t)$ falls below s^* .

As long as no player exits, the randomizations follow $\rho(t)$ and we say that the play is in *flow randomization mode*. Since beliefs change only gradually, information aggregates slowly during such a phase.

4.2 Equilibrium Path after Exits

The first exit arrives at hazard rate $n(1 - q_\theta(t))\rho(t)$. Following an exit, beliefs change drastically within the period, no matter how small Δ is. The analysis in the previous subsection implies that p must stay close to p^* as long as no player has exited. However, following an exit, the belief falls from p^* to a lower level p^- . We calculate p^- as:⁷

$$p^- = \frac{(1 - q_H)p^*}{(1 - q_H)p^* + (1 - q_L)(1 - p^*)}.$$

Consider next the equilibrium behavior after a history where $p = p^- < p^*$. There are two possibilities. First, if $\bar{p}(p, q_H, q_L, n) < p^*$, then all uninformed players exit immediately with probability one. If $\bar{p}(p, q_H, q_L, n) > p^*$, then by Remark 1, the players use an exit probability that restores p above p^* with a positive probability for the next period. Since $\Delta \downarrow 0$ means that the cost of waiting for one period, $c\Delta$, vanishes, the margin by which p should exceed p^* in the event that no player exits goes to zero. Using (8) and (9), the symmetric exit probability π that moves the belief from $p < p^*$ to p^* is given by the following:

$$(n - 1) \log \left(\frac{1 - (1 - q_H)\pi}{1 - (1 - q_L)\pi} \right) = \log \left(\frac{p^*}{1 - p^*} \right) - \log \left(\frac{p}{1 - p} \right) \equiv \tilde{p}^* - \tilde{p}.$$

When n is very large, π must be very small, and we have:

$$\log \left(\frac{1 - (1 - q_H)\pi}{1 - (1 - q_L)\pi} \right) \approx (q_H - q_L)\pi$$

and therefore, when $n \rightarrow \infty$, we have:

$$(n - 1)\pi \rightarrow \frac{\tilde{p}^* - \tilde{p}}{q_H - q_L}. \quad (31)$$

This means that the random number of players that stop within the period is distributed approximately according to a Poisson distribution with parameter $\frac{(1 - q_H)(\tilde{p}^* - \tilde{p})}{q_H - q_L}$ if $\theta = \theta_H$ and $\frac{(1 - q_L)(\tilde{p}^* - \tilde{p})}{q_H - q_L}$ if $\theta = \theta_L$.

⁷The probability of more than one exits in a single period goes to zero as $\Delta \downarrow 0$.

If k players exit, the updates in beliefs are

$$p' = \frac{p[(1 - q_H)\pi]^k [1 - (1 - q_H)\pi]^{n-1-k}}{p[(1 - q_H)\pi]^k [1 - (1 - q_H)\pi]^{n-1-k} + (1 - p)[(1 - q_L)\pi]^k [1 - (1 - q_L)\pi]^{n-1-k}},$$

$$q'_\theta = \frac{q_\theta}{1 - \pi(1 - q_\theta)}.$$

Since we are considering the limit $\Delta \downarrow 0$, updating due to private signals can be ignored.

If $k = 0$, then $p' = p^*$, and the equilibrium path goes back to the flow randomization mode described in the previous section. If k is so large that $\bar{p}(p', q'_H, q'_L, n - k) < p^*$, then the game collapses as all remaining uninformed players will exit in the next period. Otherwise, the players exit again with a relatively large probability in the next period and the exit wave continues.

The exit wave ends in one of two possible ways: If at some period $\bar{p}(p, q_H, q_L, n) < p^*$, then the game collapses and all the remaining uninformed players exit with probability one. Or, if $k = 0$ at some period, the game moves back to the flow randomization mode. As $\Delta \downarrow 0$, the succeeding periods along the exit wave are squeezed together, and the duration of the exit wave as measured in real time shrinks to zero. The duration of each flow randomization phase, on the other hand, stays strictly positive with probability one. Basically, the equilibrium path alternatives between flow randomization phases with strictly positive random duration and exit waves with negligible duration.

4.3 Discussion

We want to emphasize two properties of the symmetric equilibrium. Even as the number of players gets large, the equilibrium path displays aggregate uncertainty. The incidence and the size of the exit waves remains random, and information is only aggregated during these waves. Hence the qualitative properties of the symmetric equilibrium are quite different from the equilibria in related models such as Chamley & Gale (1994), Caplin & Leahy (1994), and Rosenberg, Solan & Vieille (2007).

It is useful to observe that even though the argument in the previous subsection proceeds by letting first $\Delta \downarrow 0$ and then letting $N \rightarrow \infty$, we would get the same results for the opposite order of limits. For a fixed Δ , we could use a Poisson approximation similar to (31) for any $s < s^*(\Delta)$ to derive the equilibrium exit distribution for large n . Letting then $\Delta \downarrow 0$ gives the same result as above.

The game always ends in a collapse. When N is large, however, a collapse reveals the state by the law of large numbers. Denote by T_θ the optimal exit time in state θ . In symmetric equilibrium with a negligible observation lag, if $\theta = \theta_H$, the collapse cannot

take place before T_H , because the last exiting player would already know the state by the time of exiting, which would contradict optimal behavior. On the other hand, the collapse occurs at latest at T_H , because no uninformed player would ever stay beyond that moment.

In contrast, if $\theta = \theta_L$, the collapse can take place at any time between $t^* > T_L$ and T_H . Denote by $\xi(t)$ the hazard rate of market collapse when $N \rightarrow \infty$, conditional on $\theta = \theta_L$. Since s is a martingale, we have:

$$s^* \lambda (1 - s^*) + (1 - p^*) \xi(t) (s_L - s^*) = 0,$$

which gives:

$$\xi(t) = \lambda \frac{s^* (1 - s^*) (s_H - s_L)}{(s^* - s_L) (s_H - s^*)}, \quad (32)$$

where we have used

$$p^* = \frac{s^* - s_L}{s_H - s_L}, \quad 1 - p^* = \frac{s_H - s^*}{s_H - s_L}.$$

Note that by (30), the hazard rate of an exit wave with $N \rightarrow \infty$, $\theta = \theta_L$, is given by $(1 - q_L(t)) \Gamma(t)$. Since not all exit waves lead to a collapse, we have $(1 - q_L(t)) \Gamma(t) > \xi(t)$.

An outside observer needs only the information about whether a collapse has taken place or not in order to update her belief on the state. If $\theta = \theta_L$, the market collapses with the hazard rate given in (32). If $\theta = \theta_H$, the game collapses at time T_H . As time goes by and the game continues, the outside observer becomes gradually more convinced that $\theta = \theta_H$. But as long as $t < T_H$, there is a positive probability that $\theta = \theta_L$.

5 Large Games

There may be other equilibria besides the symmetric and pure strategy equilibria (e.g., a subset of players behave as in a symmetric equilibrium, whereas others observe those players without revealing any information). In this section, we prove a limiting result that is valid for all equilibria of the game, when N is large and Δ is small. A large number of players allows for the possibility of almost perfect social learning. If the players were able to pool their information, their inference on the state would be accurate by the law of large numbers. In the current stopping game, the amount of information pooling is determined by equilibrium stopping behavior. Our main result shows that all equilibria of large games feature almost perfect social learning in the high state whereas social learning is slow relative to perfect information pooling in the low state.

We let $X_\theta(N, \Delta, \sigma)$ denote the (random) number of players that exit the game under equilibrium strategy profile σ when the state is θ . We will also use notation $X_\theta^t(N, \Delta, \sigma)$ to denote the (random) number of exits up to period t in state θ . Let $T_\theta(\Delta)$ be the optimal exit time of an individual player that knows the state value θ . The total probabilities of exiting the game with full information about state value are then:

$$\Xi_H(\Delta) = 1 - \alpha + \alpha(1 - \lambda\Delta)^{T_H(\Delta)}, \quad \Xi_L(\Delta) = 1 - \beta + \beta(1 - \lambda\Delta)^{T_L(\Delta)}.$$

In large games, these probabilities can be interpreted as populations shares of players that exit under full information about aggregate state (by the law of large numbers). Observe also that

$$\begin{aligned} \lim_{\Delta \rightarrow 0} T_H(\Delta) &\equiv T_H, & \lim_{\Delta \rightarrow 0} T_L(\Delta) &\equiv T_L, \\ \lim_{\Delta \rightarrow 0} \Xi_H(\Delta) &\equiv \Xi_H & \text{and} & \lim_{\Delta \rightarrow 0} \Xi_L(\Delta) \equiv \Xi_L \end{aligned}$$

are well defined.

To start the analysis, we make a preliminary observation on the maximization problem of a utilitarian social planner. Suppose that the planner does not know the state or the types of any individual players in the game but is allowed to dictate their exit strategies. By observing the realized exit decisions, the planner can update her own belief on the true state. It is easy to see that when N is large, the planner can guarantee approximately the same average payoff per player as in the case with a known state of the world. By ordering K players to exit when uninformed at T_L , the number of exits at that instant gives approximately accurate information on the state if K is large enough. If the state is revealed to be θ_L (with large probability), then all remaining uninformed players leave in the next period. If θ is revealed to be θ_H , then the remaining uninformed players leave at T_H . For Δ small enough, and N large enough, the average payoff per player is approximately optimal for each state of the world.

The main result in this section contrasts equilibrium behavior to socially optimal behavior. In equilibrium the fraction of players that exit is virtually optimal if $\theta = \theta_H$, which means that information aggregates efficiently. But if $\theta = \theta_L$, the fraction of players that exit is too low, which indicates inefficient delay:

Theorem 3 *Let σ be an equilibrium. Then:*

i) For all $\varepsilon > 0$,

$$\lim_{\Delta \rightarrow 0} \lim_{N \rightarrow \infty} \Pr\left\{\frac{X_H(N, \Delta, \sigma)}{N} - \Xi_H > \varepsilon\right\} = 0.$$

ii) There is a $\delta > 0$ such that

$$\lim_{\Delta \rightarrow 0} \lim_{N \rightarrow \infty} \Pr\left\{\Xi_L - \frac{X_L(N, \Delta, \sigma)}{N} > \delta\right\} = 1.$$

We prove Theorem 3 in the Appendix through a sequence of Lemmas. The first shows that all agents are informationally small in the sense that whenever the posterior belief on the state of the world for an outside observer with no private information is sufficiently accurate, then individual beliefs of the players are also accurate. This follows from the fact that the maximum amount of private information for any player is incorporated in her type. Accurate information thus requires information about the types of multiple players. The second Lemma shows that a large expected number of exits in any given period results in accurate beliefs for the following period. The third Lemma shows that under such circumstances, players are better off waiting for the information, as long as the period length Δ is small enough and as long as the players are not very pessimistic (i.e. their belief on the state is bounded away from zero). When Δ converges to zero, this lower bound on beliefs also converges to zero. If the state is high, then the players' beliefs stay away from zero with a high probability. Hence, in equilibrium the expected number of exits per period must be bounded. Since the total number of periods before T_H is also bounded, we get an upper bound for exits in equilibrium that is independent of the number of players and the result follows.

We emphasize the order of limits in the previous theorem. In order to maintain the discrete time nature of the model, we first fix the period length and let the number of players grow large. This allows for an arbitrarily accurate information sharing within a single period. Then we shrink the period length in order to make the cost of waiting for one period negligible.

Remark 2 *When $\beta \rightarrow 0$, all players exit eventually with probability 1 in the low state. In this case, the fraction of players that exit is ex-post optimal in both states. When $\alpha \rightarrow 1$, almost all players stay in the game forever. When $\alpha \rightarrow 1$ and $\beta \rightarrow 0$, all players exit if $\theta = \theta_L$ and almost all players stay forever if $\theta = \theta_H$. In all these cases, there is inefficient delay if $\theta = \theta_L$.*

6 Conclusion

A number of our modeling choices could be modified without affecting the qualitative nature of the results. We have assumed that information is revealed perfectly upon the arrival of a single positive signal. This assumption is made for convenience. At the expense of additional notation, we could have used a model where the different types observe positive signals at different Poisson rates. The key property that we need for our qualitative results to hold is that even the most pessimistic player wants to stay in

the game for a period of time. If some players become quickly so pessimistic that it is a dominant action for them to exit, then the state will be revealed immediately by the law of large numbers (see Rosenberg, Solan & Vieille (2007)).

The model could also be generalized beyond the two-state formulation. Even with more than two possible states, the true state would be learned at the moment of collapse if the number of players is large. Hence it can never be the case that a state is learned before it is optimal for the uninformed to exit in that given state, which means that almost all players would stay beyond the full information optimal exit time for all but the highest state. This reasoning underlines the conclusion that observational learning induces inefficient delay.

It would be straightforward to allow players to have private information on their own parameters. This would result in a symmetric equilibrium in pure cut-off strategies that would correspond to our mixed-strategy equilibrium. Letting the heterogeneity vanish in such a model would give a purification of our symmetric equilibrium.

Relaxing the assumption of irreversible actions would be more difficult. We expect the qualitative nature of our results to survive the assumption of costly re-entry, but the analysis would be much more complicated as the exit value would include the value of re-entry option, which in turn would depend on future play by other players. With fully reversible entry and exit we expect the nature of our results to change: the limit $\Delta \downarrow 0$ would allow players to easily communicate to each other their observations through an exit followed by a quick re-entry.

Finally, an interesting and challenging extension to the current model would be to allow for some form of direct payoff externalities. Then the informational aspects analyzed in this paper would combine with issues such as war of attrition, preemption, or coordination. The analytical techniques required for these cases are quite different since the value to an informed player depends on the continuation strategies of other players. We hope that our model proves useful for pursuing such extensions.

7 Appendix

7.1 Proofs of Sections 3 and 4

For the use of proofs that follow, the following lemma lists some properties of the single-observation continuation value $C_n^\pi(s, q_+, q_-)$:

Lemma 3 $C_n^\pi(s, q_+, q_-)$ is continuous in π , s , q_+ , and q_- , and:

- weakly increasing in n , with $C_1^\pi(s, q_+, q_-) = \tilde{V}_m(s)$
- weakly increasing in π , with $C_n^0(s, q_+, q_-) = \tilde{V}_m(s)$.
- strictly increasing in s , with $C_n^\pi(0, q_+, q_-) = -c\Delta$ and $C_n^\pi(1, q_+, q_-) = V^+$.
- weakly increasing in q_+ and weakly decreasing in q_- .

Whenever $C_n^\pi(s, q_+, q_-) = 0 > \tilde{V}_m(s)$, $C_n^\pi(s, q_+, q_-)$ is strictly monotonous in π , n , q_+ , and q_- .

Proof. Since $\tilde{V}_m(s)$ is continuous and strictly increasing in s , the continuity of $C_n^\pi(s, q_+, q_-)$ w.r.t. π , s , q_+ , and q_- follows directly from definitions (19) and (20). By definition (20), $C_n^\pi(s, q_+, q_-)$ is the expected value of $\tilde{V}_m(s')$, where s' is the belief updated on the basis of the signal that the single randomization event by $n - 1$ other players provides. It is clear that a signal that is superior in the sense of Blackwell increases $C_n^\pi(s, q_+, q_-)$. It is straightforward to check that the signal is indeed improved by increasing n , π , and q_+ , and by decreasing q_- . When $\tilde{V}_m(s) < 0$ and $C_n^\pi(s, q_+, q_-) = 0$, it must be that $\Pr(\tilde{V}_m(s') > 0) > 0$ and $\Pr(\tilde{V}_m(s') < 0) > 0$, which means that the optimal action (whether to exit or continue) at the next period depends crucially on the signal realization. In such a situation, the continuation value must be strictly increasing in the precision of the signal, and $C_n^\pi(s, q_+, q_-)$ is strictly increasing in n , π , and q_+ , and strictly decreasing in q_- .

The distribution of the next period s' with a given current period s , as given by (19), is stochastically first-order dominant over the same distribution with a lower current period s . Since $\tilde{V}_m(s)$ is strictly increasing in s , this means that also $C_n^\pi(s, q_+, q_-)$ is strictly increasing in s . If $s = 0$, it is optimal to stop at the next period with probability 1, so $C_n^\pi(0, q_+, q_-) = \tilde{V}_m(0) = -c\Delta$. On the other hand, $s = 1$ means that the player is good type at probability 1 and will never exit, so $C_n^\pi(1, q_+, q_-) = V^+$. ■

Proof of Lemma 1. Assume that $V_i(h^t; \sigma) < V_m(s_i(h^t; \sigma))$ at some $h^t \in H$, $i \in A(h^t)$. Then i could deviate by ignoring any information obtained by observing other

players, and replicate the behavior of a player in isolation, which would guarantee the payoff $V_m(s_i(h^t; \sigma))$. Thus, σ is not an equilibrium. In particular, exiting at $s_i(h^t, \sigma) > s^*$ would give payoff of $0 < V_m(s_i(h^t, \sigma))$. Thus, in equilibrium $s_i(h^t, \sigma) > s^*$ implies $\sigma_i(h^t) = 0$. To show that $V_i(h^t; \sigma) = V_m(s_i(h^t; \sigma))$ for at least one active player, it suffices to note that the particular player that is the next to apply a positive exit probability, must be willing to do so without any observational learning from others. Thus, this player can not have a higher payoff than a player in isolation. ■

Proof of Theorem 1. The properties of $C_n^\pi(s, q_+, q_-)$ listed in Lemma 3 imply that (22) defines a unique value $\underline{s}(n, q_+, q_-) \leq s^*$ for all relevant n , q_- , and q_+ , and (23) defines a unique value $\pi^*(n, s, q_+, q_-) \in (0, 1]$ for $s \in [\underline{s}(n, q_+, q_-), s^*]$. Also, Lemma 3 implies that $\underline{s}(n, q_+, q_-)$ and $\pi^*(n, s, q_+, q_-)$ have the following properties:

- $\underline{s}(n, q_+, q_-)$ is decreasing in n and q_+ , and increasing in q_- .
- Whenever $s \in [\underline{s}(n, q_+, q_-), s^*]$, $\pi^*(n, s, q_+, q_-)$ is strictly decreasing in n , s , and q_+ , and strictly increasing in q_- .

Let us then show that σ^S , as defined recursively in (21), is an equilibrium. We will first check whether it could be profitable for an arbitrary player i to deviate when $s(h^t, \sigma^S) < s^*$.

Denote by t_k the random number indicating the index of the k :th such period that $s(h^{t_k}, \sigma^S) < s^*$ (i.e., there are k sub-histories $h^t \subseteq h^{t_k}$, including h^{t_k} , such that $s(h^t, \sigma^S) < s^*$). Since $s(h^{t_k}, \sigma^S) < s^*$, (21) implies that $\sigma^S(h^{t_k}) > 0$. In order to check whether it could be profitable for an arbitrary player i to deviate, consider the best possible strategy for i at h^{t_k} given that other players stick to σ^S . The best-response value can be written recursively as:

$$V_i^*(h^{t_k}; \sigma_{-i}^S) = \max \left\{ \begin{aligned} &0, C_{n(h^{t_k})}^{\sigma^S(h^{t_k})}(s(h^{t_k}, \sigma^S), q_+(h^{t_k}, \sigma^S), q_-(h^{t_k}, \sigma^S)) \\ &+ \mathbb{E} \delta^{t_{k+1} - t_k} V_i^*(h^{t_{k+1}}; \sigma_{-i}^S) \end{aligned} \right\}. \quad (33)$$

Equation (33) captures the best-reponse value by decomposing it into two parts: the single-observation continuation value and the additional value that would be obtained at t_{k+1} if further observational learning makes it optimal to stay beyond that period. This additional value is the expected discounted best-response value function calculated at history $h^{t_{k+1}}$, as captured by $\mathbb{E} \delta^{t_{k+1} - t_k} V_i^*(h^{t_{k+1}}; \sigma_{-i}^S)$. The term $V_i^*(h^{t_{k+1}}; \sigma_{-i}^S)$ is again given by (33) by increasing k index by 1. This recursive formulation implies that $V_i^*(h^{t_k}; \sigma_{-i}^S)$ can be strictly positive only if there is some $m \geq k$, such that $s(h^{t_m}, \sigma^S) < s^*$ and $C_{n(h^{t_m})}^{\sigma^S(h^{t_m})}(s(h^{t_m}, \sigma^S), q_+(h^{t_m}, \sigma^S), q_-(h^{t_m}, \sigma^S)) > 0$. But equations (21), (22), and (23)

imply that this can never be the case. This means that $V_i^*(h^{t_k}; \sigma_{-1}^S) = 0$. Since σ^S gives that same payoff, i.e. $V_i(h^{t_k}; \sigma^S) = 0$, there is no profitable deviation from σ^S at h^{t_k} . Since k is arbitrary, h^{t_k} is an arbitrary history with $s(h^{t_k}, \sigma^S) < s^*$, and thus $\sigma^S(h^t)$ is a best-response whenever $s(h^t, \sigma^S) < s^*$.

On the other hand, whenever $s(h^t, \sigma^S) \geq s^*$, (21) defines $\sigma^S = 0$. Since then $V_m(s(h^t, \sigma^S)) \geq 0$, continuing must indeed be a dominant action. There is again no profitable deviation. We can now conclude that σ^S is an equilibrium.

Finally, let us confirm that σ^S is the only symmetric equilibrium. Using the monotonicity properties of the single-observation continuation value given in Lemma 3, it is straightforward to check that for every $h^t \in H$, $\sigma^S(h^t)$ as defined in (21) gives the unique symmetric randomization probability for which:

$$\begin{aligned} \text{i) } & V_i(h^t, \sigma^S) = V_m(s(h^t; \sigma^S)), \text{ and} \\ \text{ii) } & V_i(h^t, \sigma^S) \geq C_{n(h^t)}^{\sigma^S(h^t)}(s(h^t, \sigma^S), q_+(h^t, \sigma^S), q_-(h^t, \sigma^S)). \end{aligned}$$

The property i) must hold in a symmetric equilibrium by Lemma 1, and ii) must hold in any equilibrium to prevent a profitable deviation by delaying exit. This means that σ^S is the unique symmetric equilibrium. ■

Proof of Theorem 2. Denote by $\tau_k(\sigma)$ the (random) time at which σ induces the number of informative players fall to or below $N - k$ (i.e., the number of players that have revealed their private histories exceeds k at period $\tau_k(\sigma)$). Let σ^P be a profile for which (25) holds. Denote by $\sigma^{P,-i}$ the strategy profile, where i never exits, but other players stick to strategies given by σ^P (i.e. $\sigma_i^{P,-i}(h^t) = 0$ for all $h^t \in H_i$, $\sigma_j^{P,-i} = \sigma_j^P$ for all $j \neq i$). This profile is used for checking the profitability of i 's potential deviation.

Take a history $h^t \in H_i$ such that $\sigma_i^P(h^t) = 1$. The aim is to show that it is not profitable for i to deviate by choosing $\sigma_i(h^t) < 1$.

Assume that i has deviated at h^t , and due to this, is still active and uninformed at time $\tau_{N-1}(\sigma^{P,-i})$. Note that $\tau_{N-1}(\sigma^{P,-i})$ is the period at which the number of players, excluding i , that have not yet revealed their private histories goes to zero. Given that situation, consider the optimal action of i at $\tau_{N-1}(\sigma^{P,-i})$. Note that due to i 's deviation, other players' strategies are based on a false belief that i is informed. Compare the beliefs of i and j , where j is a player for which $\sigma_j^P(h^{\tau_{N-1}(\sigma^{P,-i})}) = 1$ (there must be at least one such player, because the number of informative players falls to zero exactly at period $\tau_{N-1}(\sigma^{P,-i})$). Both i and j have an identical observational history about all players $k \neq i, j$. Assuming that both are uninformed, the only source of difference in their beliefs is that while j falsely believes that i is informed at probability one, i assigns a non-trivial probability for j being informed. Assume that i were somehow able to observe

j 's information before choosing whether to continue or exit at $\tau_{N-1}(\sigma^{P,-i})$. Even then, the most favorable news that i could get would only make i 's belief identical to j 's current belief. Since by (25), single-observation continuation value is negative for j , it must also be optimal for i to exit at $\tau_{N-1}(\sigma^{P,-i})$.

Let us then consider $\tau_{N-2}(\sigma^{P,-i}) \leq \tau_{N-1}(\sigma^{P,-i})$. Assume again that i is active and uninformed at $\tau_{N-2}(\sigma^{P,-i})$. Since we have just shown that i should exit at latest at period $\tau_{N-1}(\sigma^{P,-i})$, the optimal decision of i at period $\tau_{N-2}(\sigma^{P,-i})$ is not affected by any observational learning that might take place beyond that period. This means that, by exactly the same reasoning as above, i should exit at $\tau_{N-2}(\sigma^{P,-i})$.

The same logic can now be applied step by step backwards, and we end up concluding that i should optimally stop already at the period of deviation, that is, t . So, whenever $\sigma_i^P(h^t) = 1$, there is no profitable deviation for i available.

Take then a history $h^t \in H_i$ for which $\sigma_i^P(h^t) = 0$. Then (24) and (25) imply that single-observation continuation value is positive for i , and it can not be optimal to exit. So, there is no profitable deviation for i available in that case either. Thus, σ^P is an equilibrium. ■

Proof of Lemma 2. Let $n \geq 2$ and $q_- < q_+ < 1$, and take some $\pi > 0$. Using (17) and (20), we can write the single-observation continuation payoff evaluated at s as

$$\begin{aligned} C_n^\pi(s, q_+, q_-) &= \mathbb{E} \tilde{V}_m(s') \\ &= \mathbb{E} \left[-c\Delta + s'\lambda\Delta(v + \delta V^+) + (1 - s'\lambda\Delta)\delta V_m \left(\frac{1 - \lambda\Delta}{(s')^{-1} - \lambda\Delta} \right) \right] \end{aligned} \quad (34)$$

where s' is the random value to which s jumps as a result of observing $n - 1$ other players randomize at probability π . Note that s' is independent of Δ . Since $\pi > 0$, and beliefs are martingales, there must be some $\varepsilon > 0$ such that $\Pr(s' > s + \varepsilon) > 0$.

Let $s = s^*(\Delta) - f(\Delta)$ and let $\Delta \downarrow 0$. Then the terms of the order $O(\Delta)$ in (34) go to zero, and we have:

$$\lim_{\Delta \downarrow 0} C_n^\pi(s, q_+, q_-) = \mathbb{E}[V_m(s')]. \quad (35)$$

At the same time, $\Delta \downarrow 0$ also makes $s \uparrow s^*(\Delta)$, and thus $\Pr(s' > s^*(\Delta)) > 0$. Since $V_m(s) \geq 0$ for any s and $V_m(s') > 0$ for $s' > s^*(\Delta)$, we have $\mathbb{E}[V_m(s')] > 0$, which means that $\lim_{\Delta \downarrow 0} C_n^\pi(s^*(\Delta) - f(\Delta), q_+, q_-) > 0$. Since this holds for all $\pi > 0$, definition (23) implies that $\lim_{\Delta \downarrow 0} \pi^*(n, s^*(\Delta) - f(\Delta), q_+, q_-; \Delta) = 0$. ■

7.2 Proof of Theorem 3

Lemma 4 Let $p(h^t, \sigma) = \Pr\{\theta = \theta_H | h^t\}$, i.e. the posterior probability of the good state for an outside observer with no private information. Then for all $t \leq T_H$ and $\bar{p} > 0$, there

is a $p(\bar{p}) > 0$ such that

$$p(h^t, \sigma) > 1 - p(\bar{p}) \implies p_i(h^t, \sigma) > 1 - \bar{p}$$

for all i .

Proof. From (4) and (5), we have:

$$\frac{1 - p_i(h^t, \sigma)}{1 - p(h^t, \sigma)} \leq \frac{1 - p_{i-}(h^t, \sigma)}{1 - p_{i+}(h^t, \sigma)} \leq \frac{1 - s_L^t}{s_L^t} \frac{s_H^t}{1 - s_H^t}.$$

Since s_H^t and s_L^t are bounded away from 0 and 1 for $t \in [0, T_H]$, the result follows. ■

Lemma 5 Let $X_\theta(h^t, \sigma)$ denote the random number of immediate exits at history h^t given strategy profile σ . Then

i) For all $\delta > 0$ and $\zeta > 0$ and $\underline{p} > 0$, there is a $\mu' < \infty$ such that $\mathbb{E}X_L(h^t, \sigma) > \mu'$ implies that:

$$\Pr\{p(h^{t+1}, \sigma) > 1 - \delta \mid \theta = \theta_H\} \geq 1 - \zeta,$$

whenever $p(h^t, \sigma) \geq \underline{p}$.

ii) For all $\varepsilon > 0$, there is a $K < \infty$ such that $\Pr\{X_\theta(h^t, \sigma) \geq K\mathbb{E}X_\theta(h^t, \sigma)\} \leq \varepsilon$.

Proof. i) At history h^t , the outside observer believes that player j exits with probability

$$(1 - q_\theta^j) \sigma_j \equiv \chi_\theta^j.$$

Since the randomizations are independent across agents, the mean $\mu^{X_\theta(h^t, \sigma)}$ and the variance $\text{var}^{X_\theta(h^t, \sigma)}$ in state θ are simply:

$$\mu^{X_\theta(h^t, \sigma)} = \sum_j \chi_\theta^j, \quad \text{var}^{X_\theta(h^t, \sigma)} = \sum_j \chi_\theta^j (1 - \chi_\theta^j) \leq \mu^{X_\theta(h^t, \sigma)}.$$

From (14), we know that

$$\frac{1 - q_L^j(h^t, \sigma)}{1 - q_H^j(h^t, \sigma)} = \left(\frac{1 - s_L^t \lambda \Delta}{1 - s_H^t \lambda \Delta} \right) \left(\frac{1 - q_L^j(h^{t-1}, \sigma)}{1 - q_H^j(h^{t-1}, \sigma)} \right) \frac{1 - \sigma_j(h^{t-1}) (1 - q_H^j(h^{t-1}, \sigma))}{1 - \sigma_j(h^{t-1}) (1 - q_L^j(h^{t-1}, \sigma))}.$$

Hence for all $t \geq t^*$,

$$\frac{1 - q_L^j(h^t, \sigma)}{1 - q_H^j(h^t, \sigma)} \geq 1 + \eta$$

for some $\eta > 0$. Hence we know that $\mu^{X_L(h^t, \sigma)} \geq (1 + \eta) \mu^{X_H(h^t, \sigma)}$.

Consider the random variable

$$Y_\theta(h^t, \sigma) = \frac{1}{\mu^{X_L(h^t, \sigma)}} X_\theta(h^t, \sigma).$$

Then

$$\begin{aligned}\mathbb{E}Y_H(h^t, \sigma) &\leq \frac{1}{1 + \eta} \text{ and } \mathbb{E}Y_L(h^t, \sigma) = 1, \\ \text{var}Y_H(h^t, \sigma) &= \frac{\text{var}X_H(h^t, \sigma)}{(\mu^{X_L(h^t, \sigma)})^2} \leq \frac{1}{\mu^{X_L(h^t, \sigma)}}, \\ \text{var}Y_L(h^t, \sigma) &= \frac{\text{var}X_L(h^t, \sigma)}{(\mu^{X_L(h^t, \sigma)})^2} \leq \frac{1}{\mu^{X_L(h^t, \sigma)}}.\end{aligned}$$

Hence the result follows from Chebyshev's inequality and Bayes' rule.

ii) Obvious. ■

Lemma 6 For all $\eta > 0$, there is a $\bar{p} > 0$ and a $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, and all $t < T_H$:

$$\Pr\{p(h^t, \sigma) > 1 - \bar{p} | h^{t-1}\} > \eta \implies BR_i(h^{t-1}, \sigma_{-i}) = 0,$$

where $BR_i(h^{t-1}, \sigma_{-i})$ denotes the optimal action of i at history h^{t-1} given strategy profile σ_{-i} for other players.

Proof. Since $p_{i-}(h^t, \sigma) < p(h^t, \sigma_{-i}, \sigma_i) < p_{i+}(h^t, \sigma)$, and $p_i(h^t, \sigma)$ is independent of σ_i , Lemma 4 implies that for each \hat{p} there is a $\bar{p}(\hat{p})$ such that

$$p(h^t, \sigma) > 1 - \bar{p}(\hat{p}) \implies p_i(h^t, \sigma) > 1 - \hat{p}$$

for all i and for all exit choices (also for out of equilibrium choices). Hence all i believe that

$$\Pr\{p_i(h^t, \sigma) > 1 - \hat{p}\} > \eta \text{ if } \Pr\{p(h^t, \sigma) > 1 - \bar{p}(\hat{p})\} > \eta.$$

Consider the payoff of an uninformed player i at $t < T_H$. Denote the value of a single uninformed player at history h^t when $\theta = \theta_H$ by $V_H(h^t)$. A simple continuity argument shows that for all $t < T_H$, there is a $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, $V_H(h^t) > 0$. The value function of a single player is given by $p_i(h^t) V_H(h^t) + (1 - p_i(h^t)) V_L(h^t)$ and hence the same conclusion is valid whenever p_i is large enough.

Consider now $t - 1$. If $\Pr\{p(h^t, \sigma) > 1 - \bar{p}\} > \eta$, then the cost from staying one period is $c\Delta$ and the benefit is bounded from below by $\eta V_i(h^t, \sigma)$. Hence for Δ small enough, the benefit dominates the cost and $BR_i(h^{t-1}, \sigma_{-i}) = 0$. ■

Proof of Theorem 3. i) It is clear that all uninformed players exit by $T_H(\Delta)$. Hence $\Xi_H(\Delta)$ gives a lower bound for the population share of those that exit. Consider an arbitrary equilibrium strategy profile σ . Then an upper bound for $\frac{X_H(N, \Delta, \sigma)}{N}$ is given for each $T < T_H(\Delta)$ by

$$U_H^T(N, \Delta, \sigma) = \frac{X_H^T(N, \Delta, \sigma)}{N} + \left(\frac{N - X_H^T(N, \Delta, \sigma)}{N} \right) \left(1 - \alpha + \alpha(1 - \lambda\Delta)^T \right).$$

For each $\delta > 0$, choose $T_\delta(\Delta) < T_H(\Delta)$ such that

$$\alpha(1 - \lambda\Delta)^{T_\delta(\Delta)} - \alpha(1 - \lambda\Delta)^{T_H(\Delta)} < \delta. \quad (36)$$

An obvious lower bound for $U_H^T(N, \Delta, \sigma)$ is given by $\Xi_H(\Delta)$ since all uninformed players exit by T_H .

For all $\delta > 0$, there is a \underline{p} such that $\Pr\{p(h^t, \sigma) < \underline{p} \text{ for some } h^t \mid \theta = \theta_H\} < \delta$. To see this, consider the event

$$A = \{h^t \mid p(h^t, \sigma) \leq \underline{p}\}.$$

The posterior probability of $\theta = \theta_H$ conditional on reaching A is

$$\frac{p_0 \Pr\{A \mid \theta = \theta_H\}}{p_0 \Pr\{A \mid \theta = \theta_H\} + (1 - p_0) \Pr\{A \mid \theta = \theta_L\}} \leq \underline{p}$$

by definition of the event A . Since $\Pr\{A \mid \theta = \theta_L\} \leq 1$, we have:

$$\Pr\{A \mid \theta = \theta_H\} \leq \frac{(1 - p_0) \underline{p}}{p_0 (1 - \underline{p})}.$$

Consider then paths where $p(h^t, \sigma) > \underline{p}$ for all h^t . Lemmas 5 and 6 imply that the expected number of exits is bounded by some μ in any period. The second part of Lemma 5 then implies that

$$\Pr\{X_H^{T_\delta}(\sigma, \Delta) \leq \frac{T_H - t^*}{\Delta} k(\delta') \mu(\delta')\} \geq (1 - \delta')^{\frac{T_H - t^*}{\Delta}},$$

for all $T_\delta < T_H$ for some $k(\delta') < \infty$ and for Δ small enough. By choosing δ' to be small enough $(1 - \delta')^{\frac{T_H - t^*}{\Delta}}$ can be made arbitrarily close to unity. This shows that for all $\delta > 0$, there is an \bar{N} such that for all $N \geq \bar{N}$

$$\Pr\left\{\frac{X_H^{T_\delta}(N, \Delta, \sigma)}{N} > \delta\right\} < \delta.$$

Hence for all $T_\delta < T_H$ and $\delta > 0$, there is an \bar{N} such that for all $N \geq \bar{N}$

$$\Pr\{U_H^{T_\delta}(N, \Delta, \sigma) > \delta + (1 - \delta) \left(1 - \alpha + \alpha(1 - \lambda\Delta)^{T_\delta}\right)\} < \delta.$$

Combining this with (36) yields the result.

ii) The second part follows immediately from the fact that $t^* > T_L$ and no player exits prior to t^* in any equilibrium. ■

References

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.
- Bolton, P. & C. Harris. 1999. "Strategic Experimentation." *Econometrica* 67:349–374.
- Bulow, J. & P. Klemperer. 1994. "Rational Frenzies and Crashes." *Journal of Political Economy* 102:1–23.
- Caplin, A. & J. Leahy. 1994. "Business as Usual, Market Crashes, and Wisdom After the Fact." *American Economic Review* 84:548–565.
- Chamley, C. 2004. "Delays and Equilibria with Large and Small Information in Social Learning." *European Economic Review* pp. 477–501.
- Chamley, C. & D. Gale. 1994. "Information Revelation and Strategic Delay in a Model of Investment." *Econometrica* 62:1065–1086.
- Keller, G., S. Rady & M. Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73:39–68.
- Mariotti, M. 1992. "Unused Innovations." *Economics Letters* 38:367–371.
- Rosenberg, D., E. Solan & N. Vieille. 2007. "Social Learning in One Arm Bandit Problems." *Econometrica* 75:1591–1611.
- Smith, L. & P. Sorensen. 2000. "Pathological Outcomes of Observational Learning." *Econometrica* pp. 371–398.
- Toxvaerd, F. 2008. "Strategic Merger Waves: A Theory of Musical Chairs." *Journal of Economic Theory* 140:1–26.