Learning and Information Aggregation in an Exit Game*

Pauli Murto[†] and Juuso Välimäki[‡]

This Version: February 2009

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 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". We thank Yrjö Jahnsson's foundation for financial support during the writing of this paper.

[†] 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 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 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.

---

[1] The literature on strategic experimentation considers the third case where individual experiences and actions 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.

[2] The 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.

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. In order to highlight the effects of learning and waiting, we eliminate the 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 (since the cost of waiting is small for short decision intervals). As a result, the equilibrium must be in mixed strategies. Since continuation payoffs must be positive with a positive probability (to over the instantaneous cost of delay), further exits take place with positive probability in the next period. 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 aggregated during quick bursts. We derive explicit distributions for the exit probabilities during exit waves and hazard rates for their occurrence when the number of players is

---

[3] Examples 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.

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.

Our main result shows that all equilibria of the exit game have similar information aggregation properties when the observation lag is small and the number of players is large. 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. 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.

RELATED LITERATURE

This paper is related to the literature on herding and observational learning where players have private information about a common state variable at the beginning of the game. Early papers in this literature assumed an exogenously given order of moves for the players, e.g. Banerjee (1992), Bikhchandani, Hirshleifer & Welch (1992), and Smith & Sorensen (2000). A number of later papers have endogenized the timing of action choices. Among those, the most closely related to ours is Chamley & Gale (1994).[4] In that paper a number of privately informed players consider investing in a market of uncertain aggregate profitability. The model exhibits herding with positive probability: the players' beliefs may get trapped in a region with no investment even if the market is profitable. In our model, 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, information aggregates incompletely in a single burst at the start of the game. In our model, information is revealed eventually, but at a slow rate.

Caplin & Leahy (1994) and Rosenberg, Solan & Vieille (2007) consider models with private and common values. While these papers are close to ours in their motivation, each makes a crucial modeling assumption that leads to qualitatively different information

---

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

aggregation properties to ours. Caplin and Leahy assume a continuum of players from the beginning. This assumption leads to some problems with the existence of an equilibrium and also rules out what is a key feature of our model. In our model, the actions of a large number of players result in a moderate rate of information revelation. Rosenberg, Solan & Vieille (2007) assume a finite number of players like we do, but they assume some signals that make 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 adds to 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. Section3 describes the flow of information in the game, and Section 4 provides the analysis of the symmetric and pure strategy equilibria of the model. In Section 5 we state our main result according to which all equilibria aggregate information approximately efficiently in the good state when the number of players is large. In Section 6, we characterize the symmetric equilibrium explicitly in the continuous time limit. Section 7 concludes. Appendix A contains additional details on the updating of the beliefs, Appendix B contains the proof of our main result, Theorem 3, and Appendix C contains the rest of the proofs.

2 Model

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

Before the game starts, nature chooses the (aggregate) state randomly from two alternatives: $\theta = \theta_H$ (high) and $\theta = \theta_L$ (low). Let q^0 denote the common prior $q^0 = \Pr(\theta = \theta_H)$. 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 ρ^H, while if $\theta = \theta_L$, the probability of being good is ρ^L, where $0 < \rho^L < \rho^H < 1$. The player types are drawn independently for all players. All types are initially unobservable to all players, but the parameters q^0, ρ^H, and ρ^L are common knowledge.

The information about nature's choices arrives gradually during the game as follows. In each period, each player receives a random signal $\zeta \in \{0,1\}$. Signals have two functions: they generate payoffs and transmit information. For a bad-type player, $\zeta = 0$ with probability 1. For a good player, $\Pr(\zeta = 1) = \lambda\Delta$, where λ is a commonly known parameter. Notice that informative signals arrive at a rate that depends linearly on the period length, and as a result, the real-time rate of information arrival is independent of the period length. The signal realizations across periods and players (conditional on the state and the type) are assumed to be 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, if the players knew their types, bad types would stop immediately, and good types would 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.

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. Since the information available to the uninformed players is contained in the public history of past actions, we call this simply the *history.* Formally, a history h^t in period t is defined as:

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

where $a^t = (a_1^t, ..., a_N^t)$ is the vector of individual decisions in t, and each $a_i^t \in \{0, 1\}$ is the indicator for i continuing in 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 $H_i \equiv \left\{h^t \in H \left| a_i^{t-1} = 1\right.\right\}$ the set of histories, in which i has not yet stopped. Denote by $\mathcal{N}\left(h^t\right) \equiv \left\{i \in \mathcal{N} \left| h^t \in H_i\right.\right\}$ the set of active players, and let $n\left(h^t\right)$ denote their number.

A (behavior) strategy for an uninformed player i is a mapping

$$\boldsymbol{\sigma}_i : H_i \rightarrow [0, 1]$$

that maps all histories where i is active to a stopping probability. The strategy profile is $\boldsymbol{\sigma} = (\boldsymbol{\sigma}_1, ..., \boldsymbol{\sigma}_N)$.

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 q^0. Denote by $V_i\left(h^t; \boldsymbol{\sigma}\right)$ the value of an uninformed player i after history h^t and with strategy 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 $\sigma_i\left(h^t\right)$ are best responses to $\boldsymbol{\sigma}_{-i}$ for all i and for all h^t.

3 Beliefs

3.1 Isolated player

As a useful starting point, we consider an isolated player that can only learn from her own signals. Denote by p_t the current belief of an uninformed player about her type. If the player continues for another period, and does not receive a positive signal, the new posterior p_{t+1} is obtained by Bayes' rule:

$$p_{t+1} = \frac{p_t (1 - \lambda\Delta)}{p_t (1 - \lambda\Delta) + 1 - p_t}. \tag{1}$$

Notice that apart from the prior belief at $t = 0$, the belief of a player that knows the aggregate state evolves according to the same equation (1). Denote by p^θ the player's belief on her own type conditional on state θ.

We let q denote the probability that the individual player assigns on the state being θ^H. By the law of iterated expectation,

$$p_t = q_t p_t^H + (1 - q_t) p_t^L. \tag{2}$$

Therefore

$$q_t = \frac{p_t - p_t^L}{p_t^H - p_t^L}, \tag{3}$$

and as a consequence, beliefs p_t and q_t move tightly together.

3.2 Observational learning

Exit decisions of the other players provide an additional source of information on θ and therefore also on the type of the individual player. Let A denote the (random) vector of exit decisions after history h^t. The distribution of A depends on the strategy profile $\boldsymbol{\sigma}$ of the uninformed players, current history of past exits, h^t, and the aggregate state θ (since the number of uninformed players depends on θ). Each uninformed player i must then compute her belief on the event that player j is uninformed based on h^t and $\boldsymbol{\sigma}$ conditional on state θ. In appendix A, we derive this probability and its updating formula from the knowledge of h^t and $\boldsymbol{\sigma}$. For our current purposes, it is sufficient to consider the induced random vectors $A^H(h^t, \boldsymbol{\sigma})$ and $A^L(h^t, \boldsymbol{\sigma})$ that contain the actions of all players in period t.

The total change in an individual player's beliefs within a period is the sum of the observational learning and the individual learning component as described above. It is easiest to consider the timing where an uninformed player first observes the exit decisions a_t of the other players and then receives her own signal. The updating due to observational learning is given by Bayes' rule as follows:

$$\widetilde{q}\left(h^t, a_t\right) = \frac{q(h^t) \Pr\left(A^H(h^t, \boldsymbol{\sigma}) = a^t\right)}{q(h^t) \Pr\left(A^H(h^t, \boldsymbol{\sigma}) = a^t\right) + (1 - q(h^t)) \Pr\left(A^L(h^t, \boldsymbol{\sigma}) = a^t\right)}, \tag{4}$$

where $A^{\theta}\left(h^{t}, \boldsymbol{\sigma}\right)$ is the random vector of exits after history h^{t} if state is θ and the strategy profile is given by $\boldsymbol{\sigma}$. Using this updated belief on the state, each player calculates an updated belief on her type:

$$\widetilde{p}\left(h^{t}, a^{t}\right)=\widetilde{q}\left(h^{t}, a^{t}\right) p_{t}^{H}+\left(1-\widetilde{q}\left(h^{t}, a^{t}\right)\right) p_{t}^{L}. \tag{5}$$

The individual learning induces an additional update. If the player does not see a positive signal, her final updated probability on her type is:

$$p\left(h^{t+1}\right)=\frac{\widetilde{p}\left(h^{t}, a^{t}\right)(1-\lambda \Delta)}{\widetilde{p}\left(h^{t}, a^{t}\right)(1-\lambda \Delta)+1-\widetilde{p}\left(h^{t}, a^{t}\right)}, \tag{6}$$

and therefore her belief on the state is

$$q\left(h^{t+1}\right)=\frac{p\left(h^{t+1}\right)-p_{t+1}^{L}}{p_{t+1}^{H}-p_{t+1}^{L}}.$$

4 Equilibrium analysis

4.1 Isolated player

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

$$V_{m}(p)=\max \left\{0, \widetilde{V}_{m}(p)\right\}, \tag{7}$$

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

$$\widetilde{V}_{m}(p) \equiv-c \Delta+p \lambda \Delta\left(v+\delta V^{+}\right)+(1-p \lambda \Delta) \delta V_{m}\left(\frac{p(1-\lambda \Delta)}{p(1-\lambda \Delta)+(1-p)}\right). \tag{8}$$

The optimal policy is to stop as soon as p falls below a threshold level, denoted $p^{*}(\Delta)$. The value function $V_{m}(p)$ is weakly increasing in p. By (8), $\widetilde{V}_{m}(p)$ is then strictly increasing in p. The threshold $p^{*}(\Delta)$ is obtained from (8) by setting $\widetilde{V}_{m}\left(p^{*}(\Delta)\right)=0$ and noting that $V_{m}(p)=0$ for $p<p^{*}(\Delta)$:

$$p^{*}(\Delta)=\frac{c}{\lambda\left(v+\delta V^{+}\right)}. \tag{9}$$

We shall see that $p^{*}(\Delta)$ plays a crucial role also in the model with many players. Denote by $t^*(\Delta)$ the period in which p falls below $p^*(\Delta)$ if there is no positive signal:

$$t^*(\Delta) \equiv \min\{t \in \mathbb{N}\,|p_t \leq p^*(\Delta)\}.$$

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: We denote by $p_i(h^t;\boldsymbol{\sigma})$ the belief of player i after history h^t when $\boldsymbol{\sigma}$ is the strategy profile of the players. Similarly, we let $V_i(h^t;\boldsymbol{\sigma})$ denote the continuation value of player i at history h^t when strategy profile $\boldsymbol{\sigma}$ is used.

Lemma 1. *Let $\boldsymbol{\sigma}$ be an equilibrium profile. For any $h^t \in H$,*

i) $V_i(h^t;\boldsymbol{\sigma}) \geq V_m(p_i(h^t;\boldsymbol{\sigma}))$ for all $i \in N(h^t)$,

ii) $V_i(h^t;\boldsymbol{\sigma}) = V_m(p_i(h^t;\boldsymbol{\sigma}))$ for some $i \in N(h^t)$,

iii) $\sigma_i(h^t) = 0$ whenever $p_i(h^t,\sigma) > p^*(\Delta)$.

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

It is useful to note that V_m as well as $\widetilde{V}_m$ are convex functions by usual arguments. As a result, $\mathbb{E}V_m(\widetilde{p}) > \mathbb{E}V_m(\widetilde{p}')$ and $\mathbb{E}\widetilde{V}_m(\widetilde{p}) > \mathbb{E}\widetilde{V}_m(\widetilde{p}')$ if $\mathbb{E}\widetilde{p} = \mathbb{E}\widetilde{p}'$, and $\widetilde{p}'$ second order stochastically dominates $\widetilde{p}$.

4.2 Symmetric equilibrium

In this subsection, we show that the game has a unique symmetric equilibrium. In a symmetric equilibrium, the behavior strategies of all players are identical. By Lemma 1, we know that the ex ante expected payoff of each player coincides with her value in the game as an isolated player. It also follows from Lemma 1 that players must start exiting with a positive probability immediately when $p(h^t;\boldsymbol{\sigma}) < p^*(\Delta)$. On the other hand, a pure strategy is not consistent with equilibrium if Δ is small since the informational gains from observing the others' decisions outweigh the delay costs that are proportional to Δ. Hence the symmetric equilibrium is likely to include randomized exit decisions for the uninformed players.

By symmetry, we may write $\sigma_i(h^t) = \sigma(h^t)$ for all h^t. As a result, all uninformed players have identical beliefs about the state of the world and about the informational status of other players. The key element in the Bayesian updating on the aggregate state is the exit probability of other players conditional on the state. For each state θ, we let

$\pi_i^\theta\left(h^t,\boldsymbol{\sigma}\right)=\pi^\theta\left(h^t,\boldsymbol{\sigma}\right)$ denote the probability that players other than i assign on the event that i is informed given history h^t, given the symmetric strategy profile $\boldsymbol{\sigma}$ and given the state θ. The evolution of these beliefs is also derived in Appendix A. A key observation is that for all h^t and all $\boldsymbol{\sigma}$,

$$\pi^H\left(h^t,\boldsymbol{\sigma}\right)>\pi^L\left(h^t,\boldsymbol{\sigma}\right).$$

Players other than i then believe that i exits with probability $\left(1-\pi^\theta\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma\left(h^t\right)$ if the state is θ. We use notation

$$\alpha^\theta\left(h^t,\boldsymbol{\sigma}\right)\equiv\left(1-\pi^\theta\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma\left(h^t\right)$$

to denote the individual exit probability in state θ.

We can now specialize the updating formula (4) to our symmetric setting. We calculate first the belief of an uninformed player on the aggregate state after observing k exits after a history where $n\left(h^t\right)$ players were active:

$$\widetilde{q}\left(h^t,k\right)=\frac{q\left(h^t\right)\left(\alpha^H\right)^k\left(1-\alpha^H\right)^{n\left(h^t\right)-k-1}}{q\left(h^t\right)\left(\alpha^H\right)^k\left(1-\alpha^H\right)^{n\left(h^t\right)-k-1}+\left(1-q\left(h^t\right)\right)\left(\alpha^L\right)^k\left(1-\alpha^L\right)^{n\left(h^t\right)-k-1}}. \tag{10}$$

Using the formulas (5) and (6), an uninformed player computes her belief on her own type $p\left(h^{t+1}\right)$.

For the rest of this section, it will be useful to have notation for the beliefs of an uninformed player in the next period for a given (equilibrium) exit probability $\sigma\left(h^t\right)$. We denote the (random) belief in period $t+1$ of an uninformed player following history h^t and $\sigma\left(h^t\right)$ by $p'_S\left(h^t,\sigma\left(h^t\right)\right)$.

Theorem 1. *There is a unique symmetric equilibrium $\boldsymbol{\sigma}^S$ defined by:*

$$\sigma^S\left(h^t\right)=\begin{cases}0 & \text{, if } \mathbb{E}\widetilde{V}_m\left(p'_S\left(h^t,0\right)\right)=\widetilde{V}_m\left(p'_S\left(h^t,0\right)\right)\geq 0,\\ \sigma^*\left(h^t\right) & \text{, if } \mathbb{E}\widetilde{V}_m\left(p'_S\left(h^t,0\right)\right)<0<\mathbb{E}\widetilde{V}_m\left(p'_S\left(h^t,1\right)\right),\\ 1 & \text{, if } \mathbb{E}\widetilde{V}_m\left(p'_S\left(h^t,1\right)\right)<0,\end{cases} \tag{11}$$

where $\sigma^*\left(\cdot\right)$ is the unique value implicitly defined by

$$\mathbb{E}\widetilde{V}_m\left(p'_S\left(h^t,\sigma^*\left(h^t\right)\right)\right)=0. \tag{12}$$

By Lemma 1, we know that it is never optimal for a player to exit if $p\left(h^t,\boldsymbol{\sigma}\right)>p^*\left(\Delta\right)$. The first line in (11) guarantees this. The last line implies that for low enough beliefs, all players must exit in a symmetric equilibrium. For intermediate beliefs, there is a unique (symmetric) randomization probability for other players that equates the payoff from exiting and staying for an individual player. As a result, it is a best response for each individual player to exit according to $\boldsymbol{\sigma}^*\left(h^t\right)$ in this region.

The characterization of the symmetric equilibrium yields a lower bound for $p\left(h^t, \boldsymbol{\sigma}^S\right)$ on the path where no player has exited as an immediate Corollary.

Corollary 1. *Consider the history* $h^t = (\mathbf{1}, \mathbf{1}, ..., \mathbf{1})$, *i.e. the history without any exits. Then*

$$p\left(h^t, \boldsymbol{\sigma}^S\right) \geq \frac{p^*(\Delta)(1-\lambda\Delta)}{p^*(\Delta)(1-\lambda\Delta)+1-p^*(\Delta)}.$$

4.3 Other Equilibria

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. These equilibria rely, however, on coordinated expectations on the order of moves by the players and hence they are similar to the models with exogenously given timing of actions.

A strategy profile $\boldsymbol{\sigma}$ is a pure strategy profile if $\sigma_i\left(h^t\right) \in \{0,1\}$ for all $h^t \in H$, $i \in N\left(h^t\right)$. When a pure strategy commands player i to stop, other players learn perfectly i's information. If i exits, she becomes inactive; if she stays, her type becomes public knowledge. Given a pure strategy σ, let $\widetilde{N}\left(h^t, \sigma\right) \subseteq N\left(h^t\right)$ denote the set of players, who have not yet revealed their information. We call $\widetilde{N}\left(h^t, \sigma\right)$ the set of *informative* players and $\widetilde{n}\left(h^t, \sigma\right) \leq n\left(h^t\right)$ is the number of informative players.

The information content of any single round of exit decisions at history h^t is determined by the number of informative players i such that $\sigma_i\left(h^t\right) = 1$. Denote this number by $x\left(h^t\right)$. we compute the posterior beliefs on the state conditional on the observation of k exits at history h^t as:

$$\widetilde{q}\left(h^t, k\right) = \frac{q\left(h^t\right)\left(1-\pi^H\right)^k\left(\pi^H\right)^{x\left(h^t\right)-k}}{q\left(h^t\right)\left(1-\pi^H\right)^k\left(\pi^H\right)^{x\left(h^t\right)-k}+\left(1-q\left(h^t\right)\right)\left(1-\pi^L\right)^k\left(\pi^L\right)^{x\left(h^t\right)-k}}.$$

As in the previous subsection, we let the random variable $p_P'\left(h^t, x\left(h^t\right)\right)$ denote the random belief of an uninformed player following history h^t and pure strategy exit decisions by $x\left(h^t\right)$ players. An argument that is almost identical to the previous section shows that $p_P'\left(h^t, x\left(h^t\right)\right)$ is second order stochastically decreasing in $x\left(h^t\right)$. We have the following theorem.

Theorem 2. *The exit game has a pure strategy equilibrium* $\boldsymbol{\sigma}^P$*, defined by* $x^P\left(h^t\right)$*, the number of players exiting at history* h^t*.Furthermore,* $x^P\left(h^t\right)$ *is determined by*

$$x^P\left(h^t\right)=\begin{cases}0 & \text{, if } \mathbb{E}\widetilde{V}_m\left(p'_P\left(h^t,0\right)\right)=\widetilde{V}_m\left(p'_P\left(h^t,0\right)\right)\geq 0,\\ x^*\left(h^t\right) & \text{, if } \mathbb{E}\widetilde{V}_m\left(p'_P\left(h^t,0\right)\right)<0<\mathbb{E}\widetilde{V}_m\left(p'_P\left(h^t,n\left(h^t\right)\right)\right),\\ n\left(h^t\right) & \text{, if } \mathbb{E}\widetilde{V}_m\left(p'_P\left(h^t,n\left(h^t\right)\right)\right)<0,\end{cases} \tag{13}$$

and

$$x^*\left(h^t\right)=\min\{k\,\Big|\,\mathbb{E}\widetilde{V}_m\left(p'_P\left(h^t,k\right)\right)\geq 0\}. \tag{14}$$

The pure strategy equilibrium assigns to each history a set (possibly empty) of players that reveal their private histories. Equations (13) and (14) 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 in the Appendix C 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 sufficiently delayed to discourage such deviations.

It should be noted that the game has other equilibria as well. For example those where a subset of players randomize symmetrically (within the subset) and others observe the results of these randomizations. The uninformed players not participating in the randomizations are more optimistic than the randomizing players. As a result, their continuation payoffs are positive as long as there are randomizing players left in the game.

5 Large Games

In this section, we prove a 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 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.

Let $T_\theta\left(\Delta\right)$, $\theta=H,L$, be the optimal exit period of an individual player that knows the state value θ. We use notation τ_θ to refer to the corresponding real time limit as $\Delta \to 0$:

$$\tau_\theta \equiv \lim_{\Delta \to 0} \left[T_\theta (\Delta) \cdot \Delta\right], \theta = H, L.$$

The total probability with which an individual player exits the game under full information about state value, conditional on state, is then:

$$\Xi_\theta (\Delta) = 1 - \rho^\theta + \rho^\theta (1 - \lambda\Delta)^{T_\theta(\Delta)}, \theta = H, L.$$

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). Denote

$$\Xi_\theta \equiv \lim_{\Delta \to 0} \Xi_\theta (\Delta), \theta = H, L.$$

Before stating the result, 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. When N is large and Δ is small, it is possible for the planner to guarantee approximately the same average payoff per player as in the case with a known state of the world. To see that this must indeed be the case, note that the planner can adopt the following strategy. 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 the 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.

Our main result contrasts equilibrium behavior to socially optimal behavior. Theorem 3 below states that in equilibrium, the fraction of players that exit is virtually optimal if $\theta = \theta_H$. This is equivalent to saying that virtually all uninformed players exit at the moment when they would exit if they knew the aggregate state in advance. In other words, information is aggregated efficiently. But it is important to note that this conclusion concerns only the state $\theta = H$. In contrast, if $\theta = L$, there is inefficient delay as implied by Lemma 1. The first period in which any player can exit with a positive probability, $t^* (\Delta)$, is already too late for that state. In fact, observational learning only increases delays beyond that point (this will be more concretely demonstrated in the next section). This means that observational learning amplifies the effect of aggregate uncertainty on the player's payoffs: if $\theta = H$, the players gain as observational learning improves their decisions. But if $\theta = L$, the players are hurt by additional delays induced by observational learning.

We now state the result formally. Denote by $\Omega(N, \Delta)$ the set of all equilibrium strategy profiles of the game with N players and period length Δ. Let $X_\infty(\boldsymbol{\sigma})$ denote the total number of players that exit during the game with strategy profile $\boldsymbol{\sigma}$.

Theorem 3. *For all* $\varepsilon > 0$ *and* $\delta > 0$*, there are constants* $\overline{\Delta} \in \mathbb{R}^+$ *and* $\overline{N} \in \mathbb{N}$ *such that*

$$\Pr\{\left|\frac{X_\infty(\boldsymbol{\sigma})}{N} - \Xi_H\right| > \varepsilon \,|\theta = H\} < \delta,$$

whenever $N > \overline{N}$, $\Delta < \overline{\Delta}$, *and* $\boldsymbol{\sigma} \in \Omega(N, \Delta)$.

Proof outline. We give the full proof in Appendix B using a number of Lemmas. Here we outline the main steps.

The first step in the proof shows that if N is large enough and Δ is small enough, there is an instant $\overline{t}(\Delta) < T_H(\Delta)$ such that no more than fraction $\Xi_H + \frac{\varepsilon}{2}$ of the players exit after $\overline{t}(\Delta)$ (Lemma 3). This follows from the fact that we get an upper bound for exits after $\overline{t}(\Delta)$ by assuming that no player exits prior to $\overline{t}(\Delta)$ and all uninformed exit at $\overline{t}(\Delta)$. By the law of large numbers, the share of uninformed at real time τ converges to Ξ_H as $\tau \to \tau_H$ if the true state is H. Since no player ever exits before period $t^*(\Delta)$ in equilibrium, we concentrate for the rest of the proof on periods $t \in \overline{T}$, where $\overline{T} \equiv \left\{t^*(\Delta), t^*(\Delta)+1, ..., \overline{t}(\Delta)\right\}$, and we denote the number of exits on $\overline{T}$ induced by equilibrium strategy $\boldsymbol{\sigma}$ by $\overline{X}(\boldsymbol{\sigma})$. The main task in the proof is to show that with an arbitrarily high probability, $\overline{X}(\boldsymbol{\sigma})/N < \frac{\varepsilon}{2}$.

In order to make the argument using a single belief rather than the set of beliefs of the individual players, we introduce the notion of an *outside observer*.[5] The outside observer updates her belief on the state based on the public information h^t, but in contrast to the actual players, she has no private information. We denote by $\widehat{q}(h^t, \boldsymbol{\sigma})$ the outside observer's posterior belief of $(\theta = H)$ given history h^t when the public history is generated by $\boldsymbol{\sigma}$. Lemma 4 shows that for $t \in \overline{T}$ the outsider's belief cannot diverge arbitrarily much from the belief of any uninformed player in the game. This is reasonable given that player i and the outside observer share the same belief on players $j \neq i$. For our purposes, the most important consequence of Lemma 4 is that if the outside observer is confident that the state is good at $t < \overline{t}(\Delta)$, then all players in the game must also be optimistic enough to make staying in the game a dominant action.

---

[5] For an asymmetric strategy profile $\boldsymbol{\sigma}$, the players have different beliefs on the state and on each other.

Lemma 5 uses Bayes' rule to show that conditional on $(\theta = H)$, the probability that $\Pr\{\widehat{q}(h^t, \boldsymbol{\sigma}) < \varepsilon$ for some $h^t\}$ converges to zero as $\varepsilon \to 0$. In words, this Lemma states that the outside observer cannot be arbitrarily convinced about the wrong state with positive probability.

Lemma 6 gives a bound for the numer of exits in a single period. In order to have a positive probability of many exits in a single period, $\widehat{q}(h^t, \boldsymbol{\sigma})$ would have to be very low if Δ is small. If this were not the case, then there would be a positive probability that $\widehat{q}(h^{t+1}, \boldsymbol{\sigma})$ is high, which by a corollary to Lemma 4 would mean that the continuation values of all players would be strictly positive at period $t+1$. Since the cost of waiting for one period is linear in Δ, this would mean that it is optimal for all uninformed players to stay with probability 1 at h^t if Δ is sufficiently small. On the other hand, Lemma 5 gives a lower bound on $\widehat{q}(h^t, \boldsymbol{\sigma})$, and as a result, we conclude that the number of exits cannot be very high with a positive probability at any history h^t. Lemma 7 extends this argument to the total number of exits that can take place within a time interval of random duration that is defined to end as soon as any exits take place.

Lemma 8 applies the active supermartingale Theorem by Fudenberg & Levine (1992) to show that the number of histories at which players exit is bounded. Since Fudenberg & Levine (1992) apply the theorem for the case of a fixed time interval between periods, we must come up with a slightly different argument to use it. We show that the outside observer's belief taken at the (random) periods where some players exit is an active supermartingale.

Combining Lemmas 7 and 8 then shows that for each ε, there is a finite upper bound K on the number of exits such that $\Pr\{\overline{X}(\boldsymbol{\sigma}) > K \,|\theta = H\} < \frac{\varepsilon}{2}$ whenever Δ is small enough. Since this bound is independent on N, we get the result by setting N large enough. $\square$

Remark 1. *When $\rho^H \to 1$, we have $\Xi_H \to 0$. Theorem 3 says that in the limit, the fraction of players that ever exit is negligible if $\theta = H$. When $\rho^L \to 0$, all players exit eventually with probability 1 if $\theta = L$. If $\rho^H \to 1$ and $\rho^L \to 0$, all players exit if $\theta = \theta_L$ and almost all players stay forever if $\theta = \theta_H$. In all cases, there is inefficient delay if $\theta = \theta_L$.*

6 Exit Waves

In this section, we focus on the symmetric equilibrium, and characterize it in the limit $\Delta \downarrow 0$. We have several reasons for this. First, eliminating any effects that observation lags might have on equilibrium properties is informative as such. Second, it turns out that the inherent dynamics of the model are best displayed in the limit: information aggregation happens in randomly occurring exit waves. Third, when in addition we let $N \to \infty$, we come up with a concrete illustration of Theorem 3.

It is convenient to express the limiting properties of the equilibrium in continuous time. We use variable $\tau \in [0, \infty)$ to refer to the real time in this limit. Let us stress, however, that our model is not defined in continuous time: strategies and beliefs are only defined for histories h^t, $t \in \mathbb{N}$. Actions are chosen only at real time instants $t\Delta$ for $t \in \mathbb{N}$. We extend the definitions of the variables of interest to continuous real time by the following construction. For example, we let

$$p(\tau, \Delta) \equiv p_{t(\tau, \Delta)},$$

where

$$t(\tau, \Delta) = \max\{t : t \cdot \Delta \leq \tau\},$$

and we let

$$p(\tau) \equiv \lim_{\Delta \downarrow 0} p(\tau, \Delta).$$

We denote the optimal exit threshold and time of exit for an isolated player in the limit $\Delta \downarrow 0$ by:

$$\begin{aligned} p^* &\equiv \lim_{\Delta \downarrow 0} p^* (\Delta), \\ \tau^* &\equiv \lim_{\Delta \downarrow 0} t^* (\Delta) \cdot \Delta. \end{aligned}$$

6.1 Equilibrium Path Prior to Exit

For this subsection, we assume that no player has yet exited. By Corollary 1 of Theorem 1, we know that along this history there is a function $\varepsilon(\Delta)$ satisfying

$$\lim_{\Delta \downarrow 0} \varepsilon(\Delta) = 0$$

and

$$p_t > p^* - \varepsilon(\Delta) \text{ for all } t.$$

Our first result gives an upper bound for p_t for $t \cdot \Delta \geq \tau^*$.

Lemma 2. *There is a function* $\delta(\Delta)$ *satisfying*

$$\lim_{\Delta \downarrow 0} \delta(\Delta) = 0$$

such that

$$p_t < p^* + \delta(\Delta) \text{ for all } t \geq \frac{\tau^*}{\Delta}$$

along the history where no player has exited.

This Lemma, allows us to conclude that

$$p(\tau) = p^* \text{ for all } \tau \geq \tau^*$$

along the history with no exits. Clearly then we also have

$$\dot{p}(\tau) = \lim_{d\tau \to 0} \frac{p(\tau + d\tau) - p(\tau)}{d\tau} = 0 \text{ for all } \tau \geq \tau^*. \tag{15}$$

Consider next the probability of the event that a given player i exits during a short interval $[\tau, \tau + d\tau]$ of real time. In order to interpret this probability properly, recall that we are considering the case where $\Delta \to 0$. As a result, each interval of length $d\tau$ contains many decision moments of the discrete time model and we compute the total probability of exit over all these decisions. By equations (1) and (15), it is clear that these exit probabilities must be proportional to $d\tau$ as long as there are no exits. We let $\pi^\theta(\tau)$ denote the probability that player i is informed at real time τ given that the state is θ. We let $\sigma(\tau)$ be the instantaneous exit rate for player i, i.e. an uninformed player i exits in $[\tau, \tau + d\tau]$ with probability $\sigma(\tau) d\tau$. Using the continuous time analogues of the updating formulas in Section 3.2, we can prove the following Proposition.

Proposition 1. *In the continuous time limit of the exit game, the instantaneous exit probability of an uninformed player is given by*

$$\sigma(\tau) = \lambda \frac{p^*(1-p^*)\left(p^H(\tau) - p^L(\tau)\right)}{(n-1)\left(p^* - p^L(\tau)\right)\left(p^H(\tau) - p^*\right)\left(\pi^H(\tau) - \pi^L(\tau)\right)} \text{ for } \tau \geq \tau^*. \tag{16}$$

Figure 1 shows the evolution of the belief up to the moment of the first exit. Since the beliefs evolve smoothly, and the exit probabilities are proportional to $d\tau$, we say that the game is in *flow randomization mode.*

A couple of observations on the exit rate $\sigma(\tau)$ are in order. Since exit decisions are independent, $(n-1)\sigma(\tau)\left(1 - \pi^\theta(\tau)\right)$ is the total rate at which players other than i exit.

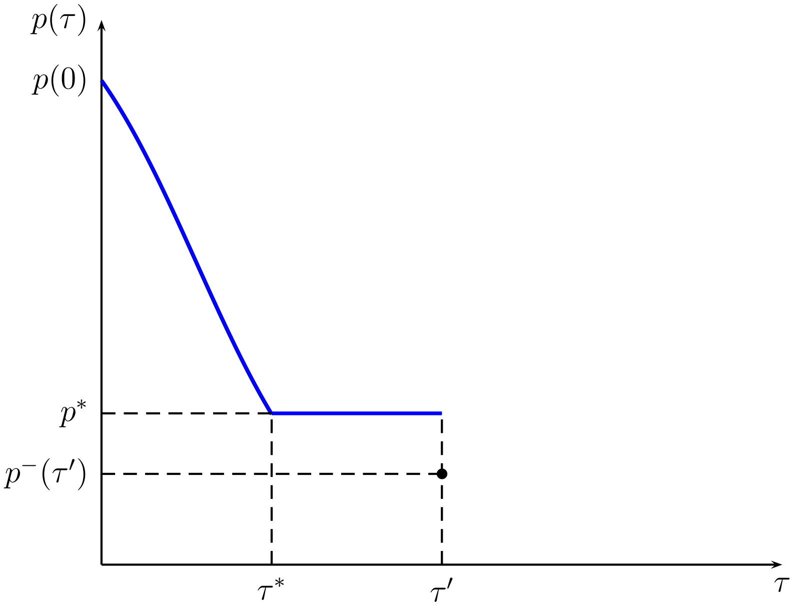

Figure 1: Beliefs in continuous time with first exit at τ'

Hence

$$(n-1)\,\sigma(\tau)\left(\pi^L(\tau)-\pi^H(\tau)\right)\equiv\Gamma(\tau)$$

is the difference in the exit rates across the two states as observed by i, and hence the effect on i's belief of state is $\Gamma(\tau)\,q(\tau)\,(1-q(\tau))$. The effect on i's belief of her own type is then

$$\Gamma(\tau)\,q(\tau)\,(1-q(\tau))\left(p^H(\tau)-p^L(\tau)\right)$$

by equation (2). On the other hand, the effect of individual learning on beliefs about her own type is given by

$$-\lambda p^*\,(1-p^*)\,.$$

The equilibrium exit rate is obtained by equating these two effects in order to keep the posterior on own type constant.

The effect of individual learning is independent of the number of players in the game. As a result, the total exit rate must also be independent of the number of players.

Since

$$q(\tau)=\frac{p^*-p^L(\tau)}{p^H(\tau)-p^L(\tau)}\leq 1,$$

we know that all uninformed players must exit at the latest at time τ_H, the moment when $p^H(\tau)=p^*$. This can also be seen in equation (16): the instantaneous exit probability of the uninformed players explodes as $\tau \to \tau_H$.

In the following subsection, we describe what happens after histories where some player has exited.

6.2 Equilibrium Path after Exits

For this subsection, it is convenient to cast the analysis in terms of the belief on the aggregate state shared by all uninformed players, q_t. We denote the continuous time limit of the symmetric equilibrium belief path prior to any exits by

$$q^* (\tau) \equiv \frac{p^* - p^L (\tau)}{p^H (\tau) - p^L (\tau)}. \tag{17}$$

The first exit arrives at hazard rate $n \left(1 - \pi^\theta (\tau)\right) \sigma (\tau)$. Following an exit, beliefs change by an amount that is large in comparison to Δ. Following an exit, the belief falls from $q^* (\tau)$ to a lower level $q^- (\tau)$. We calculate $q^- (\tau)$ as:

$$q^- (\tau) = \frac{\left(1 - \pi^H (\tau)\right) q^* (\tau)}{(1 - \pi^H (\tau))\, q^* (\tau) + (1 - \pi^L (\tau))\, (1 - q^* (\tau))}.$$

Consider next the equilibrium behavior after a history h^t where $q_t = q^- (\tau) < q^* (\tau)$ with $\tau = t \cdot \Delta$. Since we want to deal with all paths that include exits, we need to index the decision with the appropriate histories rather than just the real time instant. Recall the notation from Section 4.2 for the players' beliefs following one round of exit decisions, $p'_S (h^t, \sigma (h^t))$. By Theorem 1, there are two possibilities. If

$$\mathbb{E}\widetilde{V}_m \left(p'_S \left(h^t, 1\right)\right) < 0,$$

then all remaining players exit immediately. Otherwise, there is a $\sigma^* (h^t)$ such that

$$\mathbb{E}\widetilde{V}_m \left(p'_S \left(h^t, \sigma^* \left(h^t\right)\right)\right) = 0.$$

Notice that $\sigma^* (h^t)$ is a randomization probability in the discrete time model rather than an exit rate as in the previous subsection.

Since $\Delta \downarrow 0$ means that the cost of waiting for one period, $c \cdot \Delta$, vanishes, the margin by which $q (\tau + \Delta)$ should exceed $q^* (\tau + \Delta)$ in the event that no player exits in the current randomization goes to zero.[6] In the limit, the effect from individual learning also vanishes and we are left with the updating formula (10).

---

[6] This outcome leads to the highest possible $q (t)$ and hence also the highest $p (t)$ in the current randomization.

Proposition 2. *In the symmetric equilibrium, at aggregate state belief* $q(\tau)$ *and at real time* τ*, instantaneous randomization probability* σ *solves:*

$$(n-1)\log\left(\frac{1-\left(1-\pi^{H}(\tau)\right)\sigma}{1-\left(1-\pi^{L}(\tau)\right)\sigma}\right)=\log\left(\frac{q^{*}}{1-q^{*}}\right)-\log\left(\frac{q(\tau)}{1-q(\tau)}\right)\equiv q_{l}^{*}-q_{l}(\tau). \quad (18)$$

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

$$\log\left(\frac{1-\left(1-\pi^{H}\right)\sigma}{1-\left(1-\pi^{L}\right)\sigma}\right)\approx\left(\pi^{H}-\pi^{L}\right)\sigma,$$

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

$$(n-1)\,\sigma\left(\pi^{H}-\pi^{L}\right)\rightarrow q_{l}^{*}-q_{l}. \quad (19)$$

This means that the random number of players that stop within the period is distributed approximately according to a Poisson distribution with parameter

$$\frac{\left(1-\pi^{H}\right)\left(q_{l}^{*}-q_{l}\right)}{\left(\pi^{H}-\pi^{L}\right)}\text{ if }\theta=\theta_{H},$$

and

$$\frac{\left(1-q_{L}\right)\left(q_{l}^{*}-q_{l}\right)}{\left(\pi^{H}-\pi^{L}\right)}\text{ if }\theta=\theta_{L}.$$

When Δ is small and the number of active players n is large, the belief process during an exit wave is particularly simple to characterize. As argued above, learning from own experience is negligible when Δ is small. If n is large, each (uninformed) player exits at each moment with a small probability and as a result, conditional probability of other remaining players being informed changes very little during an exit wave. As long as the number of players that have exited is small relative to n, the beliefs of uninformed players about the aggregate state follow a (Poisson) random walk, where the current belief and the Poisson parameter for current exits is determined by the number of exits, k in the previous instant,

$$q(\tau+\Delta)=q(\tau,k).$$

If $k=0$, then $q(\tau+\Delta)=q^{*}(\tau+\Delta)$, and the equilibrium path goes back to the flow randomization mode described in the previous section (of course with fewer players than in the original game). If k is so large that uninformed players become so pessimistic that even learning the most optimistic possible news for the next period (i.e., learn that all other remaining players have received a positive signal) would not make it optimal to

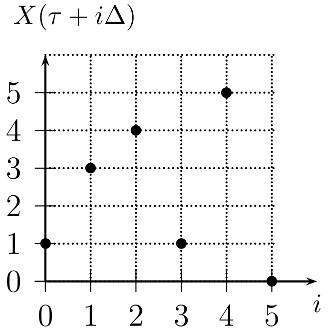

Figure 2: Exits during a wave starting at τ and ending at $\tau + 5\Delta$.

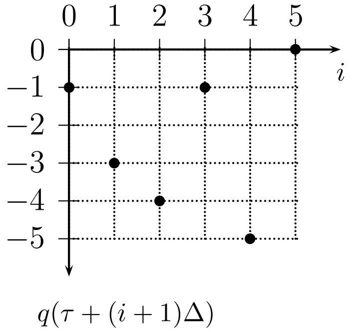

Figure 3: Beliefs during a wave starting at τ and ending at $\tau + 5\Delta$.

stay, then the game collapses as all remaining uninformed players exit at probability one. Otherwise, the players exit again with a discrete probability determined by (18), and the exit wave continues.

Hence, the exit wave ends in one of two possible ways. Either beliefs become too pessimistic and all the remaining uninformed players exit with probability one, and the game ends. Or, in some period the randomizations result in no exits. In the latter case, the game moves back to the flow randomization mode described in Section 6.1. 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.

Figures 2 and 3 show how the belief on the aggregate state is determined by the number of exits in the previous period. [7]

6.3 Discussion

Even as the number of players gets large, the equilibrium path displays aggregate uncertainty. The incidence and the size of the exit waves remain random, and information is mostly 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).

The game always ends in a collapse. When N is large, however, a collapse reveals the

---

[7] The evolution of the belief is not drawn to scale. More accurately, the figure shows the deviation of the log likelihood of the state belief as defined in Appendix A from the log likelihood at the start of the exit wave.

state by the law of large numbers. As defined in Section 5, τ_θ is 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 τ_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 τ_H, because no uninformed player would ever stay beyond that moment. This is an expression of Theorem 3: almost all players exit at the optimal time if $\theta = H$.

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

$$p^* \lambda (1 - p^*) + (1 - q^*) \xi(\tau) \left(p^L(\tau) - p^*\right) = 0,$$

which, using (17), gives:

$$\xi(\tau) = \lambda \frac{p^* (1 - p^*) \left(p^H(\tau) - p^L(\tau)\right)}{\left(p^* - p^L(\tau)\right)\left(p^H(\tau) - p^*\right)}. \tag{20}$$

Note that by the previous subsection, the hazard rate of an exit wave with $N \to \infty$, $\theta = \theta_L$, is given by

$$\begin{aligned}
\psi_L(\tau) &\equiv \lim_{n\to\infty} (n-1)\,\sigma(\tau)\left(1 - \pi^L(\tau)\right) \\
&= \lambda \frac{p^* (1 - p^*)\, p^H(\tau) - p^L(\tau)) \left(1 - \pi^L(\tau)\right)}{\left(p^* - p^L(\tau)\right)\left(p^H(\tau) - p^*\right)\left(\pi^H(\tau) - \pi^L(\tau)\right)} \\
&= \frac{\left(1 - \pi^L(\tau)\right)}{\left(\pi^H(\tau) - \pi^L(\tau)\right)} \xi(t),
\end{aligned}$$

where we have used (16) and (20). Since $\left(1 - \pi^L(\tau)\right) / \left(\pi^H(\tau) - \pi^L(\tau)\right) > 1$, we have $\psi_L(\tau) > \xi(\tau)$. The hazard rate of an exit wave is higher than that of a collapse, because not all exit waves lead to a collapse. We have plotted the hazard rates for exit waves and collapses for typical parameter values in Figure 4.

Consider for a moment an outside observer to the game, i.e. an agent that can only observe the exit decisions of the players. Except momentarily during the exit waves, a sufficient summary statistic for the outsider's belief at any point in time is whether a collapse has taken place or not. If $\theta = \theta_L$, the market collapses with the hazard rate given in (20). If $\theta = \theta_H$, the game collapses at time τ_H. As time goes by and the game continues, the outside observer becomes gradually more convinced that $\theta = \theta_H$. But as long as $\tau < \tau_H$, there is a positive probability that $\theta = \theta_L$.

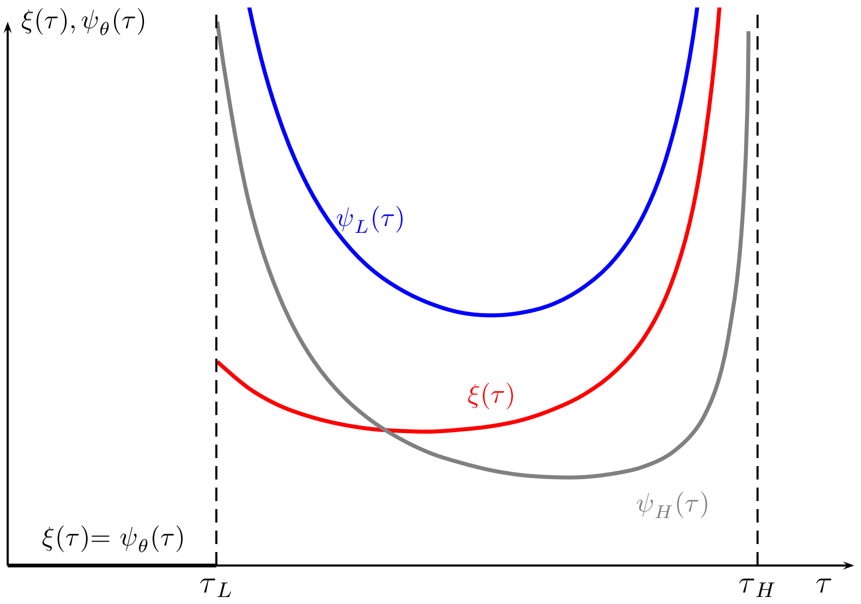

Figure 4: Hazard rates for exit waves and collapse drawn for parameter values $q_0 = 0.2, \rho^H = 0.8, \rho^L = 0.3, r = 0.05, \lambda = 0.1, v = 1, c = 0.05$.

7 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.

8 Appendix A: Beliefs

Given h^t and $\boldsymbol{\sigma}$, all beliefs can be summarized in three quantities: $q_i(h^t,\boldsymbol{\sigma})$, $\pi_j^H(h^t,\boldsymbol{\sigma})$, and $\pi_j^L(h^t,\boldsymbol{\sigma})$, $i,j \in \mathcal{N}$. Here $q_i(h^t,\boldsymbol{\sigma})$ is the belief of player i about the state (i.e. probability with which $\theta=\theta_H$), and $\pi_j^\theta(h^t,\boldsymbol{\sigma})$, $\theta\in\{\theta_H,\theta_L\}$, 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:

$$\pi_j\left(h^t,\boldsymbol{\sigma}\right)=\pi_j^H\left(h^t,\boldsymbol{\sigma}\right)q_i\left(h^t,\boldsymbol{\sigma}\right)+\pi_j^L\left(h^t,\boldsymbol{\sigma}\right)\left(1-q_i\left(h^t,\boldsymbol{\sigma}\right)\right). \tag{21}$$

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

$$p_i\left(h^t,\boldsymbol{\sigma}\right)=\Pr\{\text{type of } i \text{ is "good"} \left| h^t,\boldsymbol{\sigma}\right.\}.$$

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 $p^\theta(h^t)$ 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:

$$p^\theta\left(h^t\right)=\frac{p^\theta\left(h^0\right)\left(1-\lambda\Delta\right)^t}{p^\theta\left(h^0\right)\left(1-\lambda\Delta\right)^t+1-p^\theta\left(h^0\right)}, \tag{22}$$

where $\theta\in\{\theta_H,\theta_L\}$, and $p^H(h^0)=\rho^H$, $p^L(h^0)=\rho^L$. Using this, we may now express i's belief of her own type in terms of her belief of state:

$$\begin{aligned} p_i\left(h^t,\boldsymbol{\sigma}\right) &= q_i\left(h^t,\boldsymbol{\sigma}\right)p^H\left(h^t\right)+\left(1-q_i\left(h^t,\boldsymbol{\sigma}\right)\right)p^L\left(h^t\right) \\ &= q_i\left(h^t,\boldsymbol{\sigma}\right)\left(p^H\left(h^t\right)-p^L\left(h^t\right)\right)+p^L\left(h^t\right). \end{aligned} \tag{23}$$

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

$$q_{i+}\left(h^t,\boldsymbol{\sigma}\right) = \frac{p^H\left(h^t\right)q_i\left(h^t,\boldsymbol{\sigma}\right)}{p^H\left(h^t\right)q_i\left(h^t,\boldsymbol{\sigma}\right)+p^L\left(h^t\right)\left(1-q_i\left(h^t,\boldsymbol{\sigma}\right)\right)}, \tag{24}$$

$$q_{i-}\left(h^t,\boldsymbol{\sigma}\right) = \frac{\left(1-p^H\left(h^t\right)\right)q_i\left(h^t,\boldsymbol{\sigma}\right)}{\left(1-p^H\left(h^t\right)\right)q_i\left(h^t,\boldsymbol{\sigma}\right)+\left(1-p^L\left(h^t\right)\right)\left(1-q_i\left(h^t,\boldsymbol{\sigma}\right)\right)}. \tag{25}$$

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

$$\begin{aligned} p_{i+}^{j}\left(h^{t},\boldsymbol{\sigma}\right) &= \pi_{j}^{H}\left(h^{t},\boldsymbol{\sigma}\right) q_{i+}\left(h^{t},\boldsymbol{\sigma}\right) + \pi_{j}^{L}\left(h^{t},\boldsymbol{\sigma}\right)\left(1-q_{i+}\left(h^{t},\boldsymbol{\sigma}\right)\right), \qquad (26) \\ p_{i-}^{j}\left(h^{t},\boldsymbol{\sigma}\right) &= \pi_{j}^{H}\left(h^{t},\boldsymbol{\sigma}\right) q_{i-}\left(h^{t},\boldsymbol{\sigma}\right) + \pi_{j}^{L}\left(h^{t},\boldsymbol{\sigma}\right)\left(1-q_{i-}\left(h^{t},\boldsymbol{\sigma}\right)\right). \qquad (27) \end{aligned}$$

8.1 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

$$\widetilde{q}_{i}\left(h^{t},\boldsymbol{\sigma}\right) \equiv \log\left(\frac{q_{i}\left(h^{t},\boldsymbol{\sigma}\right)}{1-q_{i}\left(h^{t},\boldsymbol{\sigma}\right)}\right)$$

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 $\sigma_{j}\left(h^{t}\right)$. This induces the following change in i's belief:

$$\Delta_{j}\widetilde{q}_{i}\left(a_{j}\right) = \begin{cases} \log\left(\frac{1-\left(1-\pi_{j}^{H}\right)\sigma_{j}\left(h^{t}\right)}{1-\left(1-\pi_{j}^{L}\right)\sigma_{j}\left(h^{t}\right)}\right), & \text{if } a_{j}^{t}=1 \\ \log\left(\frac{\left(1-\pi_{j}^{H}\right)\sigma_{j}\left(h^{t}\right)}{\left(1-\pi_{j}^{L}\right)\sigma_{j}\left(h^{t}\right)}\right), & \text{if } a_{j}^{t}=0 \end{cases}. \qquad (28)$$

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

$$\Delta^{1}\widetilde{q}_{i}\left(a^{t}\right) = \sum_{j\in\mathcal{N}\diagdown i}\Delta_{j}\widetilde{q}_{i}\left(a_{j}^{t}\right). \qquad (29)$$

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

$$\Delta^{2}\widetilde{p} = \log\left(\frac{1-s_{H}^{t}\lambda\Delta}{1-s_{L}^{t}\lambda\Delta}\right). \qquad (30)$$

Combining the two sources of learning, we have:

$$\widetilde{p}_{i}\left(h^{t+1},\boldsymbol{\sigma}\right) = \widetilde{p}_{i}\left(h^{t},\boldsymbol{\sigma}\right) + \Delta^{1}\widetilde{p}_{i}\left(a^{t}\right) + \Delta^{2}\widetilde{p}, \qquad (31)$$

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

$$q_i\left(h^{t+1},\boldsymbol{\sigma}\right)=\frac{\exp\left(\widetilde{q}_i\left(h^{t+1},\boldsymbol{\sigma}\right)\right)}{1+\exp\left(\widetilde{q}_i\left(h^{t+1},\boldsymbol{\sigma}\right)\right)}. \tag{32}$$

Consider next the updates in a player's beliefs over another player's information, $\pi^{\theta}\left(h^t,\boldsymbol{\sigma}\right)$. 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 $\widehat{\pi}_{\theta}^{j}\left(h^t,\boldsymbol{\sigma}\right)$:

$$\widehat{\pi}_{\theta}^{j}\left(h^t,\boldsymbol{\sigma}\right)=\frac{\pi_j^{\theta}\left(h^t,\boldsymbol{\sigma}\right)}{1-\sigma_j\left(h^t\right)\left(1-\pi_j^{\theta}\left(h^t,\boldsymbol{\sigma}\right)\right)},\ \theta\in\{\theta_H,\theta_L\}. \tag{33}$$

Second, players must take into account 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:

$$\begin{aligned}\pi_j^{\theta}\left(h^{t+1},\boldsymbol{\sigma}\right) &= \widehat{\pi}_j^{\theta}\left(h^t,\boldsymbol{\sigma}\right)+\left(1-\widehat{\pi}_j^{\theta}\left(h^t,\boldsymbol{\sigma}\right)\right)p^{\theta}\left(h^t\right)\lambda\Delta\\ &= p^{\theta}\left(h^t\right)\lambda\Delta+\frac{\pi_j^{\theta}\left(h^t,\boldsymbol{\sigma}\right)\left(1-p^{\theta}\left(h^t\right)\lambda\Delta\right)}{1-\sigma_j\left(h^t\right)\left(1-\pi_j^{\theta}\left(h^t,\boldsymbol{\sigma}\right)\right)},\ \theta\in\{\theta_H,\theta_L\}. \qquad (34)\end{aligned}$$

9 Appendix B: Proof of Theorem 3

Let $\overline{\tau}$ denote the unique solution of the equation:

$$1-\rho^H\left(1-e^{-\lambda\overline{\tau}}\right)=\Xi_H+\frac{\varepsilon}{4}. \tag{35}$$

Define

$$\begin{aligned}\overline{t}\left(\Delta\right) &\equiv \max\left\{t\in\mathbb{N}\left|t\cdot\Delta<\overline{\tau}\right.\right\}, \qquad (36)\\ X^{+}\left(\boldsymbol{\sigma}\right) &\equiv n\left(h^{\overline{t}(\Delta)}\right)-n\left(h^{\infty}\right),\\ \overline{X}\left(\boldsymbol{\sigma}\right) &\equiv N-n\left(h^{\overline{t}(\Delta)}\right).\end{aligned}$$

Here $\overline{t}\left(\Delta\right)$ is the index of the last period before real time $\overline{\tau}$, $X^{+}\left(\boldsymbol{\sigma}\right)$ is the (random) number of players that exit at periods $t\geq\overline{t}\left(\Delta\right)$, and $\overline{X}\left(\boldsymbol{\sigma}\right)$ is the number of players that exit at periods $t<\overline{t}\left(\Delta\right)$. The remainder of the proof establishes an upper bound for the total number of exits. A lower bound follows easily from the fact that all uninformed players exit at latest at $T_H\left(\Delta\right)$.

Throughout the proof, we write $P_\theta(\cdot) \equiv \Pr\{\cdot\,|\theta\}$.

Lemma 3. *For all* $\varepsilon > 0$ *and* $\delta > 0$*, there is some* $\Delta' > 0$ *and* $N' < \infty$ *such that*

$$P_H\left(\frac{X^+(\boldsymbol{\sigma})}{N} > \Xi_H + \frac{\varepsilon}{2}\right) < \frac{\delta}{2}, \tag{37}$$

whenever $\boldsymbol{\sigma} \in \Omega(N,\Delta)$*,* $\Delta < \Delta'$*,* $N > N'$*.*

Proof. Consider a player that never exits the game. The probability that she is uninformed at period t given $\theta = H$ is:

$$\Xi_H^t(\Delta) = 1 - \rho^H + \rho^H(1-\lambda\Delta)^t.$$

By (35) and (36), we have then

$$\lim_{\Delta\to 0} \Xi_H^{\overline{t}(\Delta)}(\Delta) = \Xi_H + \frac{\varepsilon}{4}.$$

Assume that no player exits before time $\overline{t}(\Delta)$, and denote by $\widetilde{X}$ the number of players that are uninformed at $\overline{t}(\Delta)$. By the law of large numbers, when N is large enough, we have

$$P_H\left(\frac{\widetilde{X}}{N} > \Xi_H + \frac{\varepsilon}{2}\right) < \frac{\delta}{2}.$$

We have the result by noting that $\widetilde{X}$ is an upper bound for the number of players that may exit at periods $t \geq \overline{t}(\Delta)$ in equilibrium. □

For the rest of the proof, we concentrate on $\overline{X}(\boldsymbol{\sigma})$. As no player ever exits before period $t^*(\Delta)$ in equilibrium, we can focus on periods $t \in \overline{T}$, where $\overline{T} \equiv \{t^*(\Delta), ..., \overline{t}(\Delta)\}$. Let $\overline{H}(\Delta) \equiv \{h^t \in H \,|t \in \overline{T}\}$.

Denote by $\widehat{q}(h^t, \boldsymbol{\sigma})$ the posterior belief of $\theta = H$ at h^t, when updating is based only on public history. We call $\widehat{q}(h^t, \boldsymbol{\sigma})$ the *outside observer's belief.* We have:

Lemma 4. *For all* $\varepsilon > 0$ *there is some* $\delta > 0$ *such that*

$$\begin{aligned}
\widehat{q}\left(h^t, \boldsymbol{\sigma}\right) &> 1-\delta \Longrightarrow q_i\left(h^t, \boldsymbol{\sigma}\right) > 1-\varepsilon, \\
\widehat{q}\left(h^t, \boldsymbol{\sigma}\right) &< 1-\varepsilon \Longrightarrow q_i\left(h^t, \boldsymbol{\sigma}\right) < 1-\delta, \\
\widehat{q}\left(h^t, \boldsymbol{\sigma}\right) &< \delta \Longrightarrow q_i\left(h^t, \boldsymbol{\sigma}\right) < \varepsilon, \\
\widehat{q}\left(h^t, \boldsymbol{\sigma}\right) &> \varepsilon \Longrightarrow q_i\left(h^t, \boldsymbol{\sigma}\right) > \delta
\end{aligned}$$

for all $h^t \in \overline{H}$ *and for all* $i \in N(h^t)$*.*

Proof of Lemma 4. Denote by $q_{i-}(h^t,\boldsymbol{\sigma})$ the belief of an outside observer who observes all the other players except i. For all $h^t \in \overline{H}$, we have

$$\begin{aligned}\log\left(\frac{\widehat{q}(h^t,\boldsymbol{\sigma})}{1-\widehat{q}(h^t,\boldsymbol{\sigma})}\right) &< \log\left(\frac{q_{i-}(h^t,\boldsymbol{\sigma})}{1-q_{i-}(h^t,\boldsymbol{\sigma})}\right)+\log\left(\frac{\rho^H}{\rho^L}\right) \text{ and}\\ \log\left(\frac{q_i(h^t,\boldsymbol{\sigma})}{1-q_i(h^t,\boldsymbol{\sigma})}\right) &> \log\left(\frac{q_{i-}(h^t,\boldsymbol{\sigma})}{1-q_{i-}(h^t,\boldsymbol{\sigma})}\right)+\log\left(\frac{1-\rho^H}{1-\rho^L}\right)\end{aligned}$$

(these can be easily modified to allow $\beta = 0$ and $\alpha = 1$ utilizing the fact that $\overline{t}$ is finite).

This means that

$$\log\left(\frac{\widehat{q}(h^t,\boldsymbol{\sigma})}{1-\widehat{q}(h^t,\boldsymbol{\sigma})}\right)-\log\left(\frac{q_i(h^t,\boldsymbol{\sigma})}{1-q_i(h^t,\boldsymbol{\sigma})}\right)<\log\left(\frac{\rho^H}{\rho^L}\right)-\log\left(\frac{1-\rho^H}{1-\rho^L}\right),$$

and the result follows. ☐

Using (2) and noting that $p_t^H > p^*(\Delta)$ for all $t \le \overline{t}(\Delta)$, Lemma 4 leads to a useful corollary:

Corollary 2. *There is some* $\delta > 0$ *such that*

$$\widehat{q}\left(h^t,\boldsymbol{\sigma}\right) > 1-\delta \Longrightarrow p_i\left(h^t,\boldsymbol{\sigma}\right) > p^*(\Delta)$$

for all $h^t \in \overline{H}$ *and for all* $i \in \mathcal{N}(h^t)$.

Since $\widehat{q}(h^t,\boldsymbol{\sigma})$ is the probability assessment of a rational Bayesian agent, there is a bound for how far from truth it can be at a given probability level:

Lemma 5. *Let* $\underline{\widehat{q}}(\boldsymbol{\sigma}) \equiv \min_{t=1,\ldots,\overline{t}(\Delta)} \widehat{q}(h^t,\boldsymbol{\sigma})$. *For all* $\delta > 0$, *there is a* $q > 0$ *such that*

$$P_H\left(\underline{\widehat{q}}(\boldsymbol{\sigma}) < q\right) < \delta.$$

Proof of Lemma 5. Consider the event

$$A = \{h^t \left| \widehat{q}\left(h^t,\boldsymbol{\sigma}\right) \le q\right.\}.$$

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

$$\frac{q_0 P_H(A)}{q_0 P_H(A) + (1-q_0) P_L(A)} \le q$$

by definition of the event A. Since $P_L(A) \leq 1$, we have:

$$P_H(A) \leq \frac{(1-q_0)\,q}{q_0\,(1-q)},$$

which can be made arbitrarily small by decreasing q. $\square$

Define a sequence of stopping times $T(s)$, $s = 1, ..., S$, recursively as follows. For $s = 0$, set

$$T(0) \equiv t^*(\Delta).$$

For $s > 0$, first define $T'(s)$ as the s'th moment when at least one player exits (up to time $\bar{t}(\Delta)$):

$$T'(s) \equiv \min\left[\bar{t}(\Delta), \min\left\{t \geq T(s-1)+1 \left| n\left(h^t\right) < n\left(h^{T(s-1)}\right)\right.\right\}\right]. \tag{38}$$

To define $T(s)$, set:

$$T(s) \equiv \begin{cases} T'(s) \text{ if } P_H\left(n\left(h^{\bar{t}(\Delta)}\right) < n\left(h^{T'(s)}\right) \left| h^{T'(s)}\right.\right) > \varphi \\ \bar{t}(\Delta) \text{ otherwise} \end{cases}, \tag{39}$$

where we require that $\varphi \in \left(0, \frac{\delta}{56}\right)$. Define

$$\begin{aligned} S &\equiv \min\left\{s \left| T(s) = \bar{t}(\Delta)\right.\right\}, & (40) \\ Y_s &\equiv n\left(h^{T'(s-1)}\right) - n\left(h^{T'(s)}\right),\ s = 1, ..., S. & (41) \end{aligned}$$

We have now a finite sequence of stopping times $\{T(s)\}_{s=0}^{S}$, where S is a random integer, and where Y_s counts the number of exits between $T'(s-1)$ and $T'(s)$. Note that we use $T'(s)$ instead of $T(s)$ in the definition of Y_s. This only matters for $s = S$, and only if $T'(S) < T(S)$, in which case Y_S counts exits up to time $T'(S) < \bar{t}(\Delta)$ instead of $T(S) = \bar{t}(\Delta)$. In order to count all exits up $\bar{t}(\Delta)$, we therefore define an additional random variable

$$\overline{Y}_S \equiv n\left(h^{T'(S)}\right) - n\left(h^{T(S)}\right). \tag{42}$$

By (39), we must have:

$$P_H\left(\overline{Y}_S > 0\right) < \varphi. \tag{43}$$

The distinction between $\{T'(s)\}$ and $\{T(s)\}$ defined on the first line of (39) is made in order to ensure the following property, which is essential for Lemma 8 below to hold:

$$P_H\left(Y_s > 0 \left| h^{T(s-1)}\right.\right) > \varphi \text{ for all } s = 1, ..., S. \tag{44}$$

The total number of exits up to $\overline{t}(\Delta)$ can be written as:

$$\overline{X}(\boldsymbol{\sigma}) = \sum_{s=1}^{S} Y_s + \overline{Y}_S. \tag{45}$$

We use notation $\widehat{q}_s$, $s = 1, ..., S$ to denote the belief of the outside observer at stopping times $T(s)$:

$$\widehat{q}_s \equiv \widehat{q}\left(h^{T(s)}, \boldsymbol{\sigma}\right), \ s = 0, ..., S.$$

We also define a stochastic process L_s for $s = 1, ..., \infty$ as follows:

$$L_s \equiv \begin{cases} \frac{1-\widehat{q}_s}{\widehat{q}_s} \text{ for } s = 0, ..., S-1 \\ 0 \text{ for } s = S, ..., \infty \end{cases}. \tag{46}$$

Since L_s is the odds ratio for $\theta = L$ sampled at finite stopping times, it is well known that under $\theta = H$ this process is a supermartingale.

The key lemmas for the proof are Lemmas 6 - 8 below. However, we will first state three algebraic facts that we need in their proofs.

Claim 1. *Let* $0 < \gamma < 1$, $0 < \beta < 1$, *and* $0 < x_i < 1$ *for all* $i = 1, ..., k$. *Then*

$$\prod_{i=1}^{k}(1 - x_i) \leq 1 - \gamma \Longrightarrow \prod_{i=1}^{k}(1 - \beta x_i) \leq 1 - \beta\gamma. \tag{47}$$

Proof. Consider the maximization problem:

$$\max_{x_i > 0, i=1,...,k} \prod_{i=1}^{k}(1 - \beta x_i)$$

$$\text{subject to } \prod_{i=1}^{k}(1 - x_i) \leq 1 - \gamma.$$

It is obvious that the constraint must bind at optimum. The necessary condition for maximum is:

$$\beta \prod_{j \neq i}(1 - \beta x_j) \leq \lambda \prod_{j \neq i}(1 - \beta x_j), \ i = 1, ..., k, \tag{48}$$

where λ is the Lagrange multiplier of the constraint, and the inequality must hold as equality whenever $x_i > 0$. It is straightforward to check that (48) can hold for all i only if $x_i = x_j$ whenever $x_i x_j > 0$. It is then easy to confirm that when $0 < \beta < 1$, the maximum is attained when $x_1 > 0$, and $x_i = 0$ for $i = 2, ..., k$, and the maximum value of the problem is $1 - \beta\gamma$. $\square$

Claim 2. *There is some $\gamma > 0$ such that, for all $\Delta > 0$, $t \geq t^*(\Delta)$, and for all $\boldsymbol{\sigma}$ for which $0 < \pi_i^\theta(h^t, \boldsymbol{\sigma}) < 1$ for $\theta \in \{H, L\}$, we have:*

$$\gamma < \frac{\pi_i^L(h^t,\boldsymbol{\sigma})}{1-\pi_i^L(h^t,\boldsymbol{\sigma})} \Big/ \frac{\pi_i^H(h^t,\boldsymbol{\sigma})}{1-\pi_i^H(h^t,\boldsymbol{\sigma})} < 1 - \gamma. \tag{49}$$

Proof. Write $\kappa_\theta^t \equiv \log \frac{\pi_i^\theta(h^t,\boldsymbol{\sigma})}{1-\pi_i^\theta(h^t,\boldsymbol{\sigma})}$. Then (49) is equivalent to

$$-\log(1-\gamma) < \kappa_H^t - \kappa_L^t < -\log\gamma.$$

Writing equations (33) and (34) in the log-likelihood ratio form, yields:

$$\begin{aligned} \Delta\kappa_\theta^t &\equiv \kappa_\theta^{t+1} - \kappa_\theta^t \\ &= \log\left(\frac{1}{1-\sigma}\right) + \log\left(1 + \frac{\left(1+e^{\kappa_\theta^t}\right)p_t^\theta\lambda\Delta}{e^{\kappa_\theta^t}\left(1-p_t^\theta\lambda\Delta\right)}\right), \end{aligned} \tag{50}$$

where the first term corresponds to (33) and the second term to (34).

Let $\kappa_{H-L}^t \equiv \kappa_H^t - \kappa_L^t$ and $\Delta\kappa_{H-L}^t \equiv \kappa_{H-L}^{t+1} - \kappa_{H-L}^t = \Delta\kappa_H^t - \Delta\kappa_L^t$. Since the first term in (50) does not depend on θ, we have:

$$\Delta\kappa_{H-L}^t = \log\left(1 + \frac{\left(1+e^{\kappa_H^t}\right)p_t^H\lambda\Delta}{e^{\kappa_H^t}\left(1-p_t^H\lambda\Delta\right)}\right) - \log\left(1 + \frac{\left(1+e^{\kappa_L^t}\right)p_t^L\lambda\Delta}{e^{\kappa_L^t}\left(1-p_t^L\lambda\Delta\right)}\right). \tag{51}$$

Since $p_t^L < p_t^H$, it follows from (51) that we can have $\Delta\kappa_{H-L}^t < 0$ only when

$$\frac{e^{\kappa_H^t}}{1+e^{\kappa_H^t}} > \frac{e^{\kappa_L^t}}{1+e^{\kappa_L^t}}\frac{p_t^H}{p_t^L} > \frac{e^{\kappa_L^t}}{1+e^{\kappa_H^t}}\frac{p_t^H}{p_t^L},$$

which implies

$$\kappa_H^t - \kappa_L^t > \log\left(\frac{p_t^H}{p_t^L}\right).$$

It is easy to show that $p_{t^*(\Delta)}^H > p_{t^*(\Delta)}^L$ and that $\frac{p_t^H}{p_t^L}$ is increasing in t, and therefore $\log\left(\frac{p_t^H}{p_t^L}\right)$ is bounded from below by some $\upsilon > 0$. Thus, $\Delta\kappa_{H-L}^t$ can be negative only when $\kappa_H^t - \kappa_L^t > \upsilon$. Since $\kappa_H^{t^*(\Delta)} - \kappa_L^{t^*(\Delta)}$ is strictly positive, $\kappa_H^t - \kappa_L^t$ must be bounded from below by some positive number for all $t \geq t^* (\Delta)$. This is to say, there is some $\gamma > 0$ such that $\kappa_H^t - \kappa_L^t > -\log(1-\gamma)$ for all $h^t \in \overline{H}(\Delta)$.

On the other hand, by (51), $\Delta\kappa_{H-L}^t$ is bounded from above by

$$\Delta\kappa_{H-L}^t < \log\left(1 + \frac{\left(1+e^{\kappa_H^t}\right) p_t^H \lambda\Delta}{e^{\kappa_H^t}\left(1 - p_t^H \lambda\Delta\right)}\right) < \log(1+\varpi\Delta),$$

where ϖ is a number such that

$$\varpi > \frac{\left(1+e^{\kappa_H^t}\right) p_t^H \lambda}{e^{\kappa_H^t}\left(1 - p_t^H \lambda\Delta\right)}$$

for all $t = t^*(\Delta), ..., \overline{t}(\Delta)$. To see that such a uniform upper bound ϖ exists, note that $\lambda\Delta < 1$, $p_t^H < 1$, and $\left(1+e^{\kappa_H^t}\right)/e^{\kappa_H^t}$ is decreasing in t and bounded from above at $t = t^*(\Delta)$.

The number of periods within the real time interval $\left[t^*, \overline{T}\right]$ is bounded from above by $\frac{\overline{T}-t^*}{\Delta}$. Therefore, the total change in $\kappa_H^t - \kappa_L^t$ for $t \leq \overline{t}(\Delta)$ is bounded by:

$$\begin{aligned}\Delta\kappa_{H-L}^* &\equiv \left(\kappa_H^{\overline{t}(\Delta)} - \kappa_L^{\overline{t}(\Delta)}\right) - \left(\kappa_H^{t^*(\Delta)} - \kappa_L^{t^*(\Delta)}\right) = \sum_{t=t^*(\Delta)}^{\overline{t}(\Delta)-1} \Delta\kappa_{H-L}^t \\ &< \frac{\overline{T}-t^*}{\Delta}\log(1+\varpi\Delta).\end{aligned}$$

Since

$$\lim_{\Delta\to 0}\frac{\overline{T}-t^*}{\Delta}\log(1+\varpi\Delta) = \lim_{\Delta\to 0}\left(\overline{T}-t^*\right)\varpi\log\left[(1+\varpi\Delta)^{1/(\varpi\Delta)}\right] = \left(\overline{T}-t^*\right)\varpi,$$

it must be that $\Delta\kappa_{H-L}^*$ is bounded from above for $t \leq \overline{t}(\Delta)$, and therefore there is some $\gamma > 0$ such that $\kappa_H^t - \kappa_L^t < -\log\gamma$ for all $h^t \in \overline{H}(\Delta)$. $\square$

Claim 3. *There is some* $\gamma > 0$ *such that, for all* $\Delta > 0$, $h^t \in \overline{H}(\Delta)$, *and for all* $\boldsymbol{\sigma}$ *for which* $\pi_i^H(h^t, \boldsymbol{\sigma}) < 1$ *and* $\sigma_i(h) > 0$, *we have:*

$$\frac{\alpha_i^H(h^t,\boldsymbol{\sigma})}{\alpha_i^L(h^t,\boldsymbol{\sigma})} > \gamma, \tag{52}$$

$$\frac{1-\alpha_i^L(h^t,\boldsymbol{\sigma})}{1-\alpha_i^H(h^t,\boldsymbol{\sigma})} < 1 - \gamma\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right), \tag{53}$$

where $\alpha_i^\theta(h^t,\boldsymbol{\sigma}) \equiv \left(1-\pi_i^\theta(h^t,\boldsymbol{\sigma})\right)\sigma_i(h^t)$.

Proof. To prove (52), write:

$$\begin{aligned}\frac{\alpha_i^H\left(h^t,\boldsymbol{\sigma}\right)}{\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right)} &= \frac{\left(1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}{\left(1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)} \\ &> \frac{\left(1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)\right)\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}{\left(1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)\right)\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)} > \gamma \text{ for } \gamma \text{ small enough,}\end{aligned}$$

where the last inequality utilizes Claim 2. To prove (53), write:

$$\frac{\alpha_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\alpha_i^H\left(h^t,\boldsymbol{\sigma}\right)}\Big/\frac{\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right)}{1-\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right)} = \frac{\left(1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}{1-\left(1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}\Big/\frac{\left(1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}{1-\left(1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}$$

$$= \left(\frac{\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)} + \frac{1-\sigma_i\left(h^t\right)}{\left(1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}\right)\Big/\left(\frac{\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)} + \frac{1-\sigma_i\left(h^t\right)}{\left(1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}\right)$$

$$< \max\left(\frac{\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}\Big/\frac{\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}, \frac{1-\sigma_i\left(h^t\right)}{\left(1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}\Big/\frac{1-\sigma_i\left(h^t\right)}{\left(1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)\right)\sigma_i\left(h^t\right)}\right)$$

$$= \max\left(\frac{\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}\Big/\frac{\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}, \frac{\left(1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)\right)}{\left(1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)\right)}\right).$$

By Claim 2, we have for some $\gamma > 0$:

$$\frac{\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)}\Big/\frac{\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)} < 1-\gamma.$$

By straigthforward algebraic manipulation, this is equivalent to:

$$\frac{1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)} < 1-\gamma\pi_i^H\left(h^t,\boldsymbol{\sigma}\right). \tag{54}$$

Since $h^t \in \overline{H}\left(\Delta\right)$, $\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)$ is bounded from below by $\pi' \equiv \pi_i^H\left(h^{t^*(\Delta)},\boldsymbol{\sigma}\right) > 0$, and we have

$$\frac{1-\pi_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\pi_i^L\left(h^t,\boldsymbol{\sigma}\right)} < 1-\gamma',$$

where $\gamma' \equiv \gamma\pi' > 0$. Therefore:

$$\frac{\alpha_i^H\left(h^t,\boldsymbol{\sigma}\right)}{1-\alpha_i^H\left(h^t,\boldsymbol{\sigma}\right)}\Big/\frac{\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right)}{1-\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right)} < 1-\gamma'',$$

where $\gamma'' = \min\left(\gamma,\gamma'\right)$. Again, by algebraic manipulation, this is equivalent to:

$$\frac{1-\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right)}{1-\alpha_i^H\left(h^t,\boldsymbol{\sigma}\right)} < 1-\gamma''\alpha_i^L\left(h^t,\boldsymbol{\sigma}\right).$$

□

Lemma 6. *Let $X\left(h^{t},\boldsymbol{\sigma}\right)$ denote the random number of immediate exits at history h^{t} given strategy profile $\boldsymbol{\sigma}$. Then, for all $\delta>0$ and $\zeta>0$ and $\underline{q}>0$, there is a $\mu'<\infty$ such that $\mathbb{E}\left[X\left(h^{t},\boldsymbol{\sigma}\right)|\theta=L\right]>\mu'$ implies that:*

$$P_{H}\left(\widehat{q}\left(h^{t+1},\boldsymbol{\sigma}\right)>1-\delta\right)\geq1-\zeta,$$

whenever $\widehat{q}\left(h^{t+1},\boldsymbol{\sigma}\right)\geq\underline{q}$.

Proof. At history h^{t}, the outside observer believes that player j exits with probability $\alpha_{i}^{\theta}\left(h^{t},\boldsymbol{\sigma}\right)=\left(1-\pi_{i}^{\theta}\left(h^{t},\boldsymbol{\sigma}\right)\right)\sigma_{i}\left(h^{t}\right)$. Denote

$$\begin{aligned}
\mu_{\theta}^{X\left(h^{t},\boldsymbol{\sigma}\right)} &\equiv \mathbb{E}\left[X\left(h^{t},\boldsymbol{\sigma}\right)|\theta\right],\\
var_{\theta}^{X\left(h^{t},\boldsymbol{\sigma}\right)} &\equiv var\left[X\left(h^{t},\boldsymbol{\sigma}\right)|\theta\right]=\mathbb{E}\left[\left(X\left(h^{t},\boldsymbol{\sigma}\right)-\mu_{\theta}^{X\left(h^{t},\boldsymbol{\sigma}\right)}\right)^{2}|\theta\right].
\end{aligned}$$

Since the randomizations are independent across agents we have:

$$\mu_{\theta}^{X\left(h^{t},\boldsymbol{\sigma}\right)}=\sum_{j}\alpha_{i}^{\theta}\left(h^{t},\boldsymbol{\sigma}\right),\ var_{\theta}^{X\left(h^{t},\boldsymbol{\sigma}\right)}=\sum_{j}\alpha_{i}^{\theta}\left(h^{t},\boldsymbol{\sigma}\right)\left(1-\alpha_{i}^{\theta}\left(h^{t},\boldsymbol{\sigma}\right)\right)\leq\mu_{\theta}^{X\left(h^{t},\boldsymbol{\sigma}\right)}.$$

By (54) in Claim 3, we have:

$$\frac{1-\pi_{i}^{L}\left(h^{t},\boldsymbol{\sigma}\right)}{1-\pi_{i}^{H}\left(h^{t},\boldsymbol{\sigma}\right)}\geq1+\eta,$$

for some $\eta>0$. Consider the random variable

$$Z\left(h^{t},\sigma\right)=\frac{1}{\mu_{L}^{X\left(h^{t},\boldsymbol{\sigma}\right)}}X\left(h^{t},\boldsymbol{\sigma}\right).$$

Then

$$\begin{gathered}
\mathbb{E}\left[Z\left(h^{t},\sigma\right)|\theta=H\right]\leq\frac{1}{1+\eta},\\
\mathbb{E}\left[Z\left(h^{t},\sigma\right)|\theta=L\right]=1,\\
var\left[Z\left(h^{t},\sigma\right)|\theta=H\right]=\frac{var_{H}^{X\left(h^{t},\boldsymbol{\sigma}\right)}}{\left(\mu_{L}^{X\left(h^{t},\boldsymbol{\sigma}\right)}\right)^{2}}\leq\frac{1}{\mu_{L}^{X\left(h^{t},\boldsymbol{\sigma}\right)}},\\
var\left[Z\left(h^{t},\sigma\right)|\theta=L\right]=\frac{var_{L}^{X\left(h^{t},\boldsymbol{\sigma}\right)}}{\left(\mu_{L}^{X\left(h^{t},\boldsymbol{\sigma}\right)}\right)^{2}}\leq\frac{1}{\mu_{L}^{X\left(h^{t},\sigma\right)}}.
\end{gathered}$$

Consider the event $A = \left(Z\left(h^t,\sigma\right) \le \frac{1+\frac{1}{2}\eta}{1+\eta}\right)$. The above formulas demonstrate that

$$\lim_{\mu_L^{X\left(h^t,\sigma\right)}\to\infty} P_H\left(A\right) = 1 \text{ and } \lim_{\mu_L^{X\left(h^t,\sigma\right)}\to\infty} P_H\left(A\right) = 0.$$

Hence the claim follows from Bayes' rule. $\square$

Lemma 7. *For all* $\kappa > 0$ *and* $q > 0$*, there is some* $K < \infty$ *and* $\Delta' > 0$ *such that*

$$P_H\left(Y_s > K \left| h^{T(s-1)}\right.\right) \ge \kappa \Longrightarrow \widehat{q}\left(h^{T(s-1)}, \boldsymbol{\sigma}\right) \le q,$$

for all $s = 1, ..., S$*, and for all* $\boldsymbol{\sigma} \in \Omega\left(N, \Delta\right)$*,* $\Delta < \Delta'$*.*

Proof of Lemma 7. Fix $\kappa > 0$ and $q > 0$. Write $\mathbb{E}_\theta\left[A\right] \equiv \mathbb{E}\left[A\left|\theta\right.\right]$. Clearly, $\mathbb{E}_H\left[X\left(h^t,\boldsymbol{\sigma}\right)\right] \le \mathbb{E}_L\left[X\left(h^t,\boldsymbol{\sigma}\right)\right]$. As a first step, we will prove that there is some $K'\left(q\right) < \infty$ and $\Delta' > 0$ such that $\widehat{q}\left(h^t,\boldsymbol{\sigma}\right) > q$ implies $\mathbb{E}_H\left[X\left(h^t,\boldsymbol{\sigma}\right)\right] < K'\left(q\right)$, where $X\left(h^t;\boldsymbol{\sigma}\right)$ is the number of exits in a *single* period at any history $h^t \in \overline{H}$.

Assume that $\widehat{q}\left(h^t,\boldsymbol{\sigma}\right) > q$ and $\mathbb{E}_L\left[X\left(h^t,\boldsymbol{\sigma}\right)\right] \ge K$ for an arbitrary $K > 0$. By Lemma 6, we have:

$$\mathbb{E}_L\left[X\left(h^t,\boldsymbol{\sigma}\right)\right] \ge K \Longrightarrow P_H\left(\widehat{q}\left(h^{t+1},\boldsymbol{\sigma}\right) > 1-\phi\left(K\right)\right) > 1-\phi\left(K\right)$$

for some $\phi\left(K\right) > 0$ with $\lim\limits_{K\to\infty}\phi\left(K\right) = 0$.

By Lemma 4, $\widehat{q}\left(h^t,\boldsymbol{\sigma}\right) > q$ implies that $q_i\left(h^t,\boldsymbol{\sigma}\right) > q'$ for all i, for some $q' > 0$. On the other hand, Corollary 2 of Lemma 4 implies that with K large enough, we have:

$$\widehat{q}\left(h^{t+1},\boldsymbol{\sigma}\right) > 1-\phi\left(K\right) \Longrightarrow p_i\left(h^{t+1},\boldsymbol{\sigma}\right) > p^*\left(\Delta\right) + \gamma$$

for all i, for some $\gamma > 0$. Assume i stays at probability 1 at history h^t. The continuation payoff is then bounded from below by:

$$V_i\left(h^t,\boldsymbol{\sigma}\right) \ge -c\Delta + e^{-r\Delta}\cdot q'\left(1-\phi\left(K\right)\right)V_m\left(p^*\left(\Delta\right)+\gamma\right),$$

where $V_m\left(p^*\left(\Delta\right)+\gamma\right) > 0$ is the value of an isolated player. Clearly, $V_i\left(h^t,\boldsymbol{\sigma}\right) > 0$ if Δ is small enough and K is large enough, in which case it can not be optimal to exit, and there is thus a contradiction with the assumption $\mathbb{E}_L\left[X\left(h^t,\boldsymbol{\sigma}\right)\right] \ge K$. It follows that there is some $K'\left(q\right) < \infty$ and $\Delta' > 0$ such that

$$\widehat{q}\left(h^t,\boldsymbol{\sigma}\right) > q \Longrightarrow \mathbb{E}_L\left[X\left(h^t,\boldsymbol{\sigma}\right)\right] < K'\left(q\right) \Longrightarrow \mathbb{E}_H\left[X\left(h^t,\boldsymbol{\sigma}\right)\right] < K'\left(q\right) \tag{55}$$

for all $\boldsymbol{\sigma} \in \Omega(N, \Delta)$, $\Delta < \Delta'$.

Next, consider Y_s, the total number of exits in periods $[T'(s-1), ..., T'(s) - 1]$. Note that there is at most one period with exits, because by (38) $T'(s)$ is the period immediately after the first exit. Therefore:

$$Y_s = X\left(h^{T'(s)-1}, \boldsymbol{\sigma}\right) \leq X\left(h^{T(s-1)+k}, \boldsymbol{\sigma}\right), \tag{56}$$

where $k = \min\left(k = 0, 1, 2, ... \left| X\left(h^{T(s-1)+k}, \boldsymbol{\sigma}\right) \geq 1\right.\right)$. Since without exits outside observer's belief can not decrease, we have:

$$\widehat{q}\left(h^{T(s-1)}, \boldsymbol{\sigma}\right) > q \Longrightarrow \widehat{q}\left(h^{t'}, \boldsymbol{\sigma}\right) > q, \text{ for all } t' = T'(s-1), ..., T'(s) - 1,$$

and therefore, by (55):

$$\mathbb{E}_H\left[X\left(h^{t'}, \boldsymbol{\sigma}\right)\right] < K'(q) \text{ for all } t' = T'(s-1), ..., T'(s) - 1. \tag{57}$$

Since $X(h^t, \boldsymbol{\sigma})$ is the sum of independent Bernoully trials, the following must hold for any $t' = T'(s-1), ..., T'(s) - 1$:

$$\mathbb{E}_H\left[X\left(h^{t'}, \boldsymbol{\sigma}\right) \left| X\left(h^{t'}, \boldsymbol{\sigma}\right) \geq 1\right.\right] \leq \mathbb{E}_H\left[X\left(h^{t'}, \boldsymbol{\sigma}\right)\right] + 1 < K'(q) + 1,$$

where the second inequality uses (57). Combining this with (56), we have then:

$$\mathbb{E}_H[Y_s] < K'(q) + 1. \tag{58}$$

Since Y_s is a positive random variable, the following implication must hold:

$$P_H\left(Y_s > \frac{K'(q)+1}{\kappa} \left| h^{T(s-1)}\right.\right) > \kappa \Longrightarrow \mathbb{E}_H[Y_s] > \kappa \cdot \frac{K'(q)+1}{\kappa} = K'(q) + 1,$$

and therefore, (58) implies:

$$P_H\left(Y_s > \frac{K'(q)+1}{\kappa} \left| h^{T(s-1)}\right.\right) < \kappa.$$

Hence, setting $K > \frac{K'(q)+1}{\kappa}$, we have

$$P_H\left(Y_s > K \left| h^{T(s-1)}\right.\right) \geq \kappa \Longrightarrow \widehat{q}\left(h^{T(s-1)}, \boldsymbol{\sigma}\right) \leq q. \qquad \square$$

Lemma 8. *There is some $\gamma > 0$ such that*

$$P_H\left(\left|L_{s+1}/L_s - 1\right| > \gamma \left| h^{T(s)}\right.\right) > \gamma$$

for all $L_s > 0$.

Proof of Lemma 8. $L_s > 0$ means that $T(s) < \bar{t}(\Delta)$. The strategy of the proof is to find an event such that its probability is nonnegligible and distinctly different across states. This will ensure that the belief of an observer changes discretely at a positive probability.

Denote by $n_s \equiv n\left(h^{T(s)}\right)$ the number of active players at $T(s)$. Let us use index $m = 0, ..., M$, $M \equiv \bar{t}(\Delta) - T(s)$ to count periods after $T(s)$ up to $\bar{t}(\Delta)$:

$$t_m \equiv T(s) + m.$$

Denote by h_0^m a history at $T(s) + m$ such that no player has exited in periods $[T(s), ..., T(s) + m - 1]$:

$$h_0^m \equiv h^{T(s)} \cup \left[\bigcup_{m'=0}^{m-1} \bar{a}^{m'}\right],$$

where $\bar{a}^{m'} \equiv (1, ..., 1)$ denotes action "stay" for all active players in period $T(s) + m'$.

Without loss of generality, arrange the labels of the players players active at history h_0^m in such an order that $\alpha_1^H(h_0^m, \boldsymbol{\sigma}) \geq ... \geq \alpha_{n_s}^H(h_0^m, \boldsymbol{\sigma})$. Let $\bar{a}_k^m$ denote the action vector for period $T(s) + m$, where all k players with the highest exit probabilities stay:

$$\bar{a}_k^m \equiv \left\{ a^{T(s)+m} \left| a_i^{T(s)+m} = 1 \text{ for } i = 1, ..., k \right.\right\}$$

Define an event indexed by $m = 0, ..., M$ and $k = 1, ..., n_s$:

$$\Gamma_s(m, k) \equiv \left\{ h_0^m \cup \bar{a}_k^m \right\}.$$

$\Gamma_s(m, k)$ means that in the first m periods all the players stay, and in the $m + 1$:th period at least those k players with the highest exit probabilities stay. We have:

$$\begin{aligned} P_H\left(\Gamma_s(m,k) \left| h^{T(s)}\right.\right) &= \prod_{m'=0}^{m-1} \prod_{i=1}^{n_s} \left(1 - \alpha_i^H\left(h_0^{m'}\right)\right) \cdot \prod_{i=1}^{k} \left(1 - \alpha_i^H\left(h_0^m\right)\right), \\ P_L\left(\Gamma_s(m,k) \left| h^{T(s)}\right.\right) &= \prod_{m'=0}^{m-1} \prod_{i=1}^{n_s} \left(1 - \alpha_i^L\left(h_0^{m'}\right)\right) \cdot \prod_{i=1}^{k} \left(1 - \alpha_i^L\left(h_0^m\right)\right). \end{aligned}$$

Let us now choose $m = m^*$ and $k = k^*$ in such a way that

$$\varphi < P_H\left(\Gamma_s\left(m^*, k^*\right)\left|h^{T(s)}\right.\right) < 1 - \varphi. \tag{59}$$

To see that it is always possible to choose such m^* and k^*, note that $P_H\left(\Gamma_s\left(m, k\right)\left|h^{T(s)}\right.\right)$ is decreasing in m and k, and by (44), we have

$$P_H\left(\Gamma_s\left(M, n_s\right)\left|h^{T(s)}\right.\right) = P_H\left(Y_{s+1} = 0\left|h^{T(s)}\right.\right) < 1 - \varphi,$$

and on the other hand:

$$P_H\left(\Gamma_s\left(0, 1\right)\left|h^{T(s)}\right.\right) \geq \pi_1^H\left(h^{T(s)}, \boldsymbol{\sigma}\right) > \pi_i^H\left(h^{t^*(\Delta)}, \boldsymbol{\sigma}\right).$$

We have:

$$\frac{P_L\left(\Gamma_s\left(m^*, k^*\right)\left|h^{T(s)}\right.\right)}{P_H\left(\Gamma_s\left(m^*, k^*\right)\left|h^{T(s)}\right.\right)} = \frac{\prod_{m'=0}^{m^*-1}\prod_{i=1}^{k^*}\left(1 - \alpha_i^L\left(h_0^{m'}\right)\right) \cdot \prod_{i=1}^{k^*}\left(1 - \alpha_i^L\left(h_0^{m^*}\right)\right)}{\prod_{m'=0}^{m^*-1}\prod_{i=1}^{k^*}\left(1 - \alpha_i^H\left(h_0^{m'}\right)\right) \cdot \prod_{i=1}^{k^*}\left(1 - \alpha_i^H\left(h_0^{m}\right)\right)}$$

$$= \prod_{m'=0}^{m^*-1}\prod_{i=1}^{k^*}\frac{\left(1 - \alpha_i^L\left(h_0^{m'}\right)\right)}{\left(1 - \alpha_i^H\left(h_0^{m'}\right)\right)} \cdot \prod_{i=1}^{k^*}\frac{\left(1 - \alpha_i^L\left(h_0^{m}\right)\right)}{\left(1 - \alpha_i^H\left(h_0^{m}\right)\right)}.$$

Applying (53) in Claim 3 to all terms, we have:

$$\frac{P_L\left(\Gamma_s\left(m^*, k^*\right)\left|h^{T(s)}\right.\right)}{P_H\left(\Gamma_s\left(m^*, k^*\right)\left|h^{T(s)}\right.\right)} < \prod_{m'=0}^{m^*-1}\prod_{i=1}^{k^*}\left(1 - \gamma'\alpha_i^L\left(h_0^{m'}\right)\right) \cdot \prod_{i=1}^{k^*}\left(1 - \gamma'\alpha_i^L\left(h_0^{m}\right)\right)$$

for some $\gamma' > 0$. Noting that

$$\begin{aligned}\prod_{m'=0}^{m^*-1}\prod_{i=1}^{k^*}\left(1 - \alpha_i^L\left(h_0^{m'}\right)\right) \cdot \prod_{i=1}^{k^*}\left(1 - \alpha_i^L\left(h_0^{m}\right)\right) &= P_L\left(\Gamma_s\left(m, k\right)\left|h^{T(s)}\right.\right) = \\ &< P_H\left(\Gamma_s\left(m, k\right)\left|h^{T(s)}\right.\right) < 1 - \varphi,\end{aligned}$$

we can use Claim 1 to conclude:

$$\frac{P_L\left(\Gamma_s\left(m^*, k^*\right)\left|h^{T(s)}\right.\right)}{P_H\left(\Gamma_s\left(m^*, k^*\right)\left|h^{T(s)}\right.\right)} < 1 - \gamma'\varphi. \tag{60}$$

By Bayes' rule, (60) means that

$$\frac{1-\widehat{q}_{s+1}}{\widehat{q}_{s+1}}\left|\Gamma_s\left(m^*,k^*\right)\right. < \left(1-\gamma'\varphi\right)\left(\frac{1-\widehat{q}_s}{\widehat{q}_s}\right),$$

which implies that $L_{s+1}/L_s - 1 < -\gamma'\varphi$. Since by (59), $\Gamma_s\left(m^*,k^*\right)$ is an event that occurs at a strictly positive probability, there must then be some $\gamma > 0$ such that

$$P_H\left(\left|L_{s+1}/L_s - 1\right| > \gamma \left|h^{T(s)}\right.\right) > \gamma. \qquad \square$$

Lemma 8 says that L_s, $s \in \mathbb{N}$ is an *active supermartingale* with activity γ, as defined in Fudenberg & Levine (1992). We need this property to apply Theorem A.1. of Fudenberg & Levine (1992), which we restate here for convenience:

Theorem 4 (Fudenberg and Levine)**.** *Let* $l_0 > 0$*,* $\varepsilon > 0$*, and* $\gamma \in (0,1)$ *be given. For each* $\underline{L}$*,* $0 < \underline{L} < l_0$*, there is some* $K < \infty$ *such that*

$$\Pr\left(\sup_{k>K} L_k \leq \underline{L}\right) \geq 1-\varepsilon$$

for every active supermartingale L *with* $L_0 = l_0$ *and activity* γ*.*

To apply this Theorem in our context, note that by (46):

$$L_k \leq \underline{L} \Longleftrightarrow \left(\widehat{q}_k \geq \frac{1}{1+\underline{L}} \text{ or } k \geq S\right).$$

By Corollary 2 of Lemma (4), we can therefore set $\underline{L}$ in such a way that

$$L_k \leq \underline{L} \Longrightarrow \left(p_i\left(h^{T(k)},\boldsymbol{\sigma}\right) > p^*\left(\Delta\right) \forall i \text{ or } k \geq S\right). \tag{61}$$

Now we apply Lemma 8 and Theorem 4, and set K so high that

$$P_H\left(\sup_{k>K} L_k \leq \underline{L}\right) \geq 1-\varphi. \tag{62}$$

Since by Lemma 1 no player is ever willing to exit if $p_i\left(h^t,\boldsymbol{\sigma}\right) > p^*\left(\Delta\right)$, equations (39), (61), and (62) imply:

$$P_H\left(S > K\right) \leq \varphi. \tag{63}$$

Applying Lemma 5, choose $q > 0$ such that

$$P_H\left(\underline{\widehat{q}}\left(\boldsymbol{\sigma}\right) < q\right) < \varphi.$$

Then, applying Lemma 7, set $K' < \infty$ high enough and $\Delta' > 0$ low enough to guarantee:

$$\widehat{q}\left(h^{T(s)}, \boldsymbol{\sigma}\right) > q \Longrightarrow P_H\left(Y_s > K' \left| h^{T(s)}\right.\right) < \frac{\varphi}{K}.$$

Then:

$$P_H\left(\max_{s=1,\dots,S} Y_s > K' \left| S \leq K, \underline{\widehat{q}}(\boldsymbol{\sigma}) \geq q\right.\right) < K \cdot \frac{\varphi}{K} = \varphi,$$

whenever $\boldsymbol{\sigma} \in \Omega(N, \Delta)$, $\Delta \in (0, \Delta')$. This means that:

$$\begin{aligned} P_H\left(\max_{s=1,\dots,S} Y_s > K'\right) &\leq P_H\left(\underline{\widehat{q}}(\boldsymbol{\sigma}) < q\right) + P_H(S > K) \\ &\quad + \left[1 - P_H\left(\underline{\widehat{q}}(\boldsymbol{\sigma}) < q\right) - P_H(S > K)\right] \cdot \varphi \\ &< \varphi + \varphi + \varphi = 3\varphi. \end{aligned}$$

Since $\overline{X}(\boldsymbol{\sigma}) = \sum_{s=1}^{S} Y_s + \overline{Y}_S$, and $P_H\left(\overline{Y}_S = 0\right) > 1 - \varphi$ by (43), we have:

$$\begin{aligned} P_H\left(\overline{X}(\boldsymbol{\sigma}) < K \cdot K'\right) &= P_H\left(\sum_{s=1}^{S} Y_s + \overline{Y}_S < K \cdot K'\right) \qquad (64) \\ &> P_H(S < K) \cdot P_H\left(\max_{s=1,\dots,S} Y_s < K'\right) \cdot P_H\left(\overline{Y}_S = 0\right) \\ &> (1 - \varphi)(1 - 3\varphi)(1 - \varphi) > (1 - 3\varphi)^3 > 1 - 27\varphi \\ &> 1 - \frac{\delta}{2}, \end{aligned}$$

where the last inequality follows from the fact $\varphi \in \left(0, \frac{\delta}{56}\right)$. Equation (64) holds for all $\boldsymbol{\sigma} \in \Omega(N, \Delta)$ as long as $\Delta < \Delta'$. There is no restriction on N, and both K and K' are independent on N. The total number of exits is

$$X_\infty(\boldsymbol{\sigma}) = \overline{X}(\boldsymbol{\sigma}) + X^+(\boldsymbol{\sigma}). \qquad (65)$$

Combining (37), (64), and (65), we have $P_H\{\frac{X_\infty(\boldsymbol{\sigma})}{N} > \Xi_H + \varepsilon\} < \delta$ if N is sufficiently large.

10 Appendix C: Other Proofs

Proof of Lemma 1. Let σ_i^m be the optimal strategy of the isolated player. $V_i(p; \boldsymbol{\sigma}) \geq V_m(p)$ since σ_i^m is a feasible strategy in the exit game. Let

$$i^* \in \{i \in \mathcal{N} \left| \left(\min_t \sigma_i\left(h^t\right) > 0\right) \leq \left(\min_t \sigma_j\left(h^t\right) > 0\right) \text{ for all } j \right.\}.$$

Then $V_{i^*}(p;\boldsymbol{\sigma}) \leq V_m(p)$ since no other player exits before i^* with positive probability. $\square$

Proof of Theorem 1. It is clear that no player exits if the payoff from staying for one period is positive. If $\widetilde{V}_m(p'_S(h^t,0)) < 0$, then the players must exit in equilibrium with positive probability. We start by showing that $\mathbb{E}\widetilde{V}_m(p'_S(h^t,\sigma(h^t)))$ is strictly increasing in $\sigma(h^t)$ if $\mathbb{E}\widetilde{V}_m(p'_S(h^t,1)) > 1$.

First we observe that $\widetilde{q}(h^t,\cdot)$ as defined in (10) is (strictly) second order stochastically decreasing and continuous in $\sigma(h^t)$. To see this, notice first that the vector of exit decisions is independent across players conditional on the state. Given the observation on the action of a single player, we can construct other players' belief on the aggregate state (suppressing the arguments):

$$\begin{aligned}\widetilde{q}\left(h^t,0\right) &= \frac{q\left(1-\alpha^H\right)}{q\left(1-\alpha^H\right)+\left(1-q\right)\left(1-\alpha^L\right)} = \frac{1}{1+\frac{1-q}{q}\frac{1-(1-\pi^L)\sigma}{1-(1-\pi^H)\sigma}},\\ \widetilde{q}\left(h^t,1\right) &= \frac{qa^H}{qa^H+\left(1-q\right)a^L} = \frac{1}{1+\frac{1-q}{q}\frac{1-\pi^L}{1-\pi^H}}.\end{aligned}$$

Since $\widetilde{q}(h^t,0)$ is increasing in σ and $\widetilde{q}(h^t,1)$ is independent of σ and since $\mathbb{E}\widetilde{q}(h^t,\cdot) = q$ for all σ by the law of iterated expectation, the posterior after observing a single exit decision is second order stochastically decreasing in σ. By independence (conditional on the state), the claim extends to the full vector of exit decisions.

By equation (5), $\widetilde{p}(h^t,\cdot)$ is linear in $\widetilde{q}(h^t,\cdot)$. By equation (6), $p(h^{t+1})$ is convex in $\widetilde{p}(h^t,\cdot)$. Finally, $\mathbb{E}\widetilde{V}_m(p'_S(h^t,\sigma(h^t)))$ is convex in p', and as a result,

$$\mathbb{E}\widetilde{V}_m\left(p'_S\left(h^t,\sigma\left(h^t\right)\right)\right)$$

is increasing in $\sigma(h^t)$.

If

$$\Pr\{\widetilde{V}_m\left(p'_S\left(h^t,1\right)\right) = 0\} > 0,$$

and

$$\Pr\{\widetilde{V}_m\left(p'_S\left(h^t,1\right)\right) > 0\} > 0,$$

then $\mathbb{E}\widetilde{V}_m(p'_S(h^t,\sigma(h^t)))$ is strictly increasing in $\sigma(h^t)$ for all values of $\sigma(h^t)$ for which $\Pr\{\widetilde{V}_m(p'_S(h^t,\sigma(h^t))) > 0\} > 0$. $\square$

Proof of Corollary 1. Note that $\widetilde{q}(h^t,0) \geq \widetilde{q}(h^t,k)$ for all k and therefore $p\left(\left(h^t,\mathbf{1}\right),\boldsymbol{\sigma}^S\right) \geq$

$p\left(\left(h^t, a^t\right), \boldsymbol{\sigma}^S\right)$ for all a^t. As a result,

$$\mathbb{E}\widetilde{V}_m\left(p'\left(h^t, \sigma^*\left(h^t\right)\right)\right) = 0$$

implies that

$$\widetilde{V}_m\left(p\left(\left(h^t, \mathbf{1}\right), \boldsymbol{\sigma}^S\right)\right) > 0,$$

and therefore

$$p\left(\left(h^t, \mathbf{1}\right), \boldsymbol{\sigma}^S\right) > p^*\left(\Delta\right). \qquad \square$$

Proof of Theorem 2. Denote by $t_k\left(\boldsymbol{\sigma}\right)$ the (random) time at which $\boldsymbol{\sigma}$ 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 $t_k\left(\boldsymbol{\sigma}\right)$). Let $\boldsymbol{\sigma}^P$ be a profile for which (13) - (14) hold. Denote by $\boldsymbol{\sigma}^{P,-i}$ the strategy profile, where i never exits, but other players use strategies given by $\boldsymbol{\sigma}^P$ (i.e. $\sigma_i^{P,-i}\left(h^t\right) = 0$ for all $h^t \in H_i$, $\boldsymbol{\sigma}_j^{P,-i} = \boldsymbol{\sigma}_j^P$ 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 $\sigma_i^P\left(h^t\right) = 1$. The aim is to show that it is not profitable for i to deviate by choosing $\sigma_i\left(h^t\right) < 1$.

Assume that i has deviated at h^t, and due to this, is still active and uninformed at time $t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$. Note that $t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$ 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 $t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$. 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 $\sigma_j^P\left(h^{t_{N-1}\left(\sigma^{P,-i}\right)}\right) = 1$ (there must be at least one such player, because the number of informative players falls to zero exactly at period $t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$). 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 probability less than one for j being informed. Assume that i is given an option to observe j's information before choosing whether to continue or exit at $t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$. Even then, the most favorable news that i could get would only make i's belief identical to j's current belief. Since it is optimal for j to exit, it must also be optimal for i to exit at $t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$.

Let us then consider $t_{N-2}\left(\boldsymbol{\sigma}^{P,-i}\right) \leq t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$. Assume again that i is active and uninformed at $t_{N-2}\left(\boldsymbol{\sigma}^{P,-i}\right)$. Since we have just shown that i should exit at latest at period $t_{N-1}\left(\boldsymbol{\sigma}^{P,-i}\right)$, the optimal decision of i at period $t_{N-2}\left(\boldsymbol{\sigma}^{P,-i}\right)$ 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 $t_{N-2}\left(\boldsymbol{\sigma}^{P,-i}\right)$.

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 $\sigma_i^P\left(h^t\right)=1$, there is no profitable deviation for i available.

Take then a history $h^t \in H_i$ for which $\sigma_i^P\left(h^t\right)=0$. Then (13) - (14) imply that one-period 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, $\boldsymbol{\sigma}^P$ is an equilibrium. $\square$

Proof of Lemma 2. Suppose not. Then there exists a $\delta>0$, a sequence $\{\Delta_k\}\to 0$, and a corresponding sequence of real time instants $\{t\left(\Delta_k\right)\}$ along the path with no exits such that

$$p\left(t\left(\Delta_k\right)\right)\geq p^*\left(\Delta\right)+\delta.$$

For each such Δ_k, let

$$\widehat{t}\left(\Delta_k\right)=\{\min t\geq\frac{\tau^*}{\Delta}\left|p\left(t\right)\geq p^*\left(\Delta\right)+\delta\right.\}.$$

Since $V_m\left(p\right)$ is strictly increasing in p for $p\geq p^*\left(\Delta\right)$, there is an $\varepsilon>0$ such that $V_m\left(p+\delta\right)>\varepsilon$. There is always a positive probability ζ that no player exits in a given period (since they might all be informed). Therefore the payoff from staying is bounded from below by

$$e^{-r\Delta_k}\zeta\varepsilon-c\Delta_k.$$

For small enough Δ_k, this is positive. By the definition of $\widehat{t}\left(\Delta_k\right)$, and the characterization of symmetric equilibria in the previous section, we have derived a contradiction. $\square$

Proof of Proposition 1. Write

$$p\left(\tau\right)=q\left(\tau\right)p^H\left(\tau\right)+\left(1-q\left(\tau\right)\right)p^L\left(\tau\right). \tag{66}$$

Along the symmetric equilibrium path, for $\tau\geq\tau^*$ we must have:

$$\dot{p}\left(\tau\right)=\dot{p}^L\left(\tau\right)+\dot{q}\left(\tau\right)\left(p^H\left(\tau\right)-p^L\left(\tau\right)\right)+q\left(\tau\right)\left(\dot{p}^H\left(\tau\right)-\dot{p}^L\left(\tau\right)\right)=0. \tag{67}$$

Using (1), we have along the history with no exits:

$$\dot{p}^\theta\left(\tau\right)=-\lambda p^\theta\left(\tau\right)\left(1-p^\theta\left(\tau\right)\right)\text{ for }\theta\in\{H,L\}. \tag{68}$$

Substituting from (68) and (66) into (67) gives:

$$\dot{q}(\tau) = -\lambda q(\tau)(1-q(\tau))\left(p^H(\tau) - p^L(\tau)\right) + \frac{p^*(1-p^*)}{p^H(\tau) - p^L(\tau)}. \tag{69}$$

On the other hand, given some $\sigma(\tau)$, the differential change in q induced by the observational learning from other players can be calculated as:

$$q + dq =$$
$$\frac{q\left(1-p^H+p^H(1-\lambda dt)\right)\left(1-\alpha^H dt\right)^{n-1}}{q\left(1-p^H+p^H(1-\lambda dt)\right)\left(1-\alpha^H dt\right)^{n-1} + (1-q)\left(1-p^L+p^L(1-\lambda dt)\right)\left(1-\alpha^L dt\right)^{n-1}},$$

where

$$\alpha^\theta(\tau) = \left(1-\pi^\theta\right)\sigma(\tau). \tag{70}$$

Letting $d\tau \to 0$ reduces this to:

$$\dot{q}(\tau) = q(\tau)(1-q(\tau))\left(-\lambda p^H(\tau) - (n-1)\alpha^H + \lambda p^L(\tau) + (n-1)\alpha^L\right). \tag{71}$$

Equating (69) and (71), and using (66) and (70), we get

$$\sigma(\tau) = \frac{\lambda p^*(1-p^*)\left(p^H(\tau) - p^L(\tau)\right)}{(n-1)\left(p^* - p^L(\tau)\right)\left(p^H(\tau) - p^*\right)\left(\pi^H(\tau) - \pi^L(\tau)\right)}. \qquad \square$$

Proof of Proposition 2. The formula in 18 follows from 10 by requiring that:

$$\widetilde{q}\left(h^t, 0\right) = q^*. \qquad \square$$

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.

Fudenberg, D. & D. Levine. 1992. “Maintaining a Reputation When Strategies are Imperfectly Observed.” *Review of Economic Studies* 59:561–579.

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.