

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

Pauli Murto[†] and Juuso Välimäki[‡]

This Version: April 2010

Abstract

We analyze information aggregation in a stopping game with uncertain payoffs that are correlated across players. 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

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 a timing game with pure informational externalities. We show

^{*}We would like to thank numerous seminar audiences and, in particular Dirk Bergemann, Hikmet Gunay, Johannes Hörner, Godfrey Keller, Elan Pavlov, Sven Rady, Larry Samuelson, 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.

[†]Aalto University, Department of Economics, Helsinki, Finland, pauli.murto@hse.fi

[‡]Aalto University, Department of Economics, Helsinki, Finland, juuso.valimaki@hse.fi

that even though private information accumulates steadily over time in our model, it is revealed 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. Demand may be high or low. For a population of monopolistically competing firms, the market is profitable to a larger fraction of firms if demand is high. Learning from others is useful to the extent that it can be used to determine the overall demand level. It is not sufficient, however, as it may be that the product of an individual firm does not appeal to the consumers even when demand is high.

We use a standard discounted single-player experimentation model in discrete time to represent private learning. Players do not know their type at the beginning of the game. Over time, they learn by observing signals that are correlated with their true payoff type. We assume binary types. Good types gain in expected terms by staying in the game while 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.¹ Uninformed players become more pessimistic as time passes and the optimal strategy is to exit the game once a threshold level of pessimism is reached.

When many players face the same decision problem and their types are correlated, then observations of other players' behavior becomes another source of information to the players. We model correlation across players by assuming that there is a binary state of the world that determines the probability distribution of individual types. In the high state, a higher fraction of the players are of the good type. Conditional on the state, the players' types are identically and independently distributed. Whenever the exit decisions of a given player are sensitive to her information, her actions reveal information about her informational type and hence also about the state of the world (since more players are informed in the high state). 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 prevents the players from waiting 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.

¹The actual form of information revelation is not very important for the logic of our model. The important assumption is that it takes some time for any player to become so pessimistic that exiting is optimal.

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 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.

When a player exits, the beliefs of the other players become more pessimistic. This means that from the perspective of an individual uninformed player, it is suddenly optimal to exit as well. However, immediate exit by all uninformed players would release so much information that an individual player might find it optimal to wait (since the cost of delay is small for frequent periods). As a result, the equilibrium must be in mixed strategies that balance the incentives to exit and wait. If there are further exits from the market as a result of the randomization, pessimism persists and yet another round of randomizations is called for. We call this phase of consecutive exits an *exit wave*.² Only when there is a period with no exits, a sufficient level of optimism is restored in the market and the exit wave ends. An exit wave thus ends either in a collapse of the game where the last uninformed player exits, or in a reversion to the flow mode. In the symmetric equilibrium, play fluctuates randomly between these two modes until a collapse ends the game. In Section 7, we argue that exit decisions are bunched together also in similar models where we allow for heterogeneity between the uninformed players.³

When the number of players is increased towards infinity, the pooled information on the aggregate state becomes accurate. One might conjecture that conditional on the state aggregate randomness would vanish 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 modes 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 random bursts. We compute the exit probabilities during exit waves and the hazard rate for their occurrence when the number of players is large.

We show that even though in the large game limit the aggregate state must eventually be revealed, this can never happen before the efficient stopping time conditional on the

²Other 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, whereas our waves are purely informational.

³Without observational learning, exit decisions would then vary as a function of the private information of the players.

true state. In particular, this means that information is aggregated efficiently in the high state: almost all uninformed players exit as if they knew the true state. But if the state is low, information aggregation fails: the players learn the state too late. In terms of the payoffs, the possibility to observe other players' behavior helps the good types by making them less likely to exit the game. At the same time, observational learning hurts the bad types as they are encouraged to stay in the game too long.

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 gradual private learning about common values. While these papers are close to ours in their motivation, each makes a crucial modeling assumption that leads to qualitatively different information aggregation properties to ours. Caplin and Leahy assume a continuum of players from the beginning. This assumption 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 signals that may make a player so pessimistic after one period that exiting is the dominant strategy right away. As a result, when the number of players is increased towards infinity, 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 &

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

Leahy (1994) and Rosenberg, Solan & Vieille (2007).⁵ Our model displays a qualitatively different form of endogenous information transmission between the players. In our model, private information is accumulated gradually and this information is revealed in randomly occurring exit waves. In Section 7, we discuss some extensions of our main model where the pattern of information aggregation remains similar to the main model.

Another difference to the literature mentioned above is that 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). This makes it possible to have a richer analysis of the welfare consequences of social learning. The pure common values case is obtained in our model as a special case.

Our paper is also related to the literature on strategic experimentation. That literature focuses on the private provision of public information rather than aggregation of privately held information. Examples of such models are Bolton & Harris (1999) and Keller, Rady & Cripps (2005). The key difference is that in those models the signals of all players are publicly observable, whereas in our model the players see only each other's actions.

The paper is organized as follows. Section 2 sets up the discrete time model and Subsection 2.1 presents an alternative interpretation for the model as a model of irreversible investment. Section 3 describes the flow of information in the game, and Section 4 provides the analysis of the symmetric equilibrium. In Section 5, we discuss information aggregation in large games. In Section 6, we characterize the symmetric equilibrium in the continuous time limit. Section 7 presents a number of extensions of the basic model and some suggestions for future work. Appendix A contains the proofs for all results given in Sections 3, 4, and 6. The proofs for Section 5 are longer and they are given in Appendix B.

2 Model

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

Before the game starts, nature chooses an (aggregate) state randomly from two alternatives: $\theta \in \{H, L\}$. Let q^0 denote the common prior $q^0 = \Pr\{\theta = H\}$. After choosing the state, nature chooses randomly and independently the type of each player. Each

⁵The main emphasis in Rosenberg, Solan & Vieille (2007) is on the characterization of the symmetric equilibrium in games with a small number of players.

player is either good or bad, but the players do not know their type at the beginning of the game. If $\theta = H$, the probability of being good is ρ^H , while if $\theta = L$, the probability of being good is ρ^L , where $0 \leq \rho^L < \rho^H \leq 1$. In the special case, where $\rho^H = 1$ and $\rho^L = 0$, the players' types are perfectly correlated and the game is one of pure common values. Conditional on the state, 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 the aggregate state and the individual types arrives gradually during the game. As long as a player stays in the game, she receives a random signal $\zeta \in \{0, 1\}$ in each period. 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 $\lambda > 0$ is a commonly known parameter. Upon receiving the signal, the players collect a payoff $\zeta \cdot v$, where v is the value of the good signal $\zeta = 1$ for the players. Notice that informative signals arrive at a rate that depends linearly on the period length, and as a result, information arrives at a constant rate that is independent of period length. The signal realizations across periods and players (conditional on the state and the type) are assumed to be independent. 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 good signal are informed, other players are uninformed.

At the beginning of each period t , all active players i make a binary decision a_i^t . They either exit, $a_i^t = 0$, or continue, $a_i^t = 1$. Exiting is costless, but irreversible: once a player exits, she becomes inactive and receives the outside option payoff 0. Hence we require that whenever $a_i^t = 0$, then $a_i^s = 0$ for all $s > t$. We call player i active in period t if she has stayed in the game up to that point in time. We denote by \mathcal{N} the set of active players and we let n denote their number. If the player continues in the game, she pays the (opportunity) cost $c \cdot \Delta$. The cost c and the benefit v are parameters for which we assume $0 < c < \lambda v$. We also assume risk neutrality (i.e. we measure the costs and benefits in utils). The expected payoff per period is $(\lambda v - c) \Delta > 0$ for a good player and $-c\Delta < 0$ for a bad player. This means that if the players knew their types, bad types would exit immediately, and good types would never exit.⁶

Within each period, the players act simultaneously, but they know each others' previous actions. However, they do not observe each others' signals, and therefore they do not know whether the other players are informed or uninformed.

The history of player i consists of her private history of own past signals, and the

⁶In Section 7, we discuss the possibility that the value of the signal is different in the two states.

public history consisting of the actions of all players. Since a good signal reveals fully the player's type, the uninformed have never observed the good signal. Conditional on a good signal, it is a dominant strategy to stay in the game forever. It is therefore sufficient to specify the exit behavior of the uninformed players. For the uninformed players, all relevant information is contained in the public history of past actions, and therefore we call this public information simply the *history*. Formally, a history h^t in period t is a sequence of actions:

$$h^t = (a^0, a^1, \dots, a^{t-1}),$$

where $a^t = (a_1^t, \dots, a_N^t)$. Denote by H^t the set of all such histories up to t and let $H = \bigcup_{t=0}^{\infty} H^t$. A history $h^\infty = \{a^t\}_{t=0}^{\infty}$ gives a sequence of action profiles for the entire game.

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

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

that maps all histories where i is active to an exit probability. A strategy profile in the game is a vector $\sigma = (\sigma_1, \dots, \sigma_N)$.

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

2.1 Interpretation as an Investment Game

We can interpret the game as an investment model where a number of firms have the option of undertaking an irreversible investment. The project qualities are correlated across firms. A good project yields $c\Delta > 0$ per period whereas a bad project yields $(c - \lambda v)\Delta < 0$ per period. The fixed investment cost is normalized to zero. Before undertaking the project, each firm learns about the quality of her individual potential project as follows. In each period of this prior experimentation phase, a "failure" signal occurs with probability $\lambda\Delta$ if and only if the project is bad.⁷ The firms do not, however, earn the proceeds $c\Delta$ before the actual investment has been made.

To see that this is equivalent to our exit game, consider the capitalized value of undertaking the action "exit" in our original model. If the player type is bad, then by exiting the player avoids the fixed cost $c\Delta$ from today to eternity. Therefore the

⁷See Décamps & Mariotti (2004) for another investment model with this kind of learning.

capitalized value of exit is equal to the value of investing in a "good" project in the investment model. A "good" type in the original model avoids the cost $c\Delta$ by exiting, but at the same time she forgoes the expected payoff $\lambda v\Delta$ per period. The net capitalized value of exit is then equal to the value of investing in a "bad" project. This shows that the two models are isomorphic (with "good" types interpreted as "bad" projects and "bad" types as "good" projects).

3 Beliefs

In this section we fix a strategy profile σ and describe the two different forms of learning in our model. First, as long as a player stays in the game, she receives in every period a direct signal on her type. The strength of this signal is exogenously given, and Bayesian updating resulting from such signals has been studied extensively in the literature. Subsection 3.1 below describes this private learning.

The second form of learning depends on the publicly observed actions and is endogenous in our model. Since all past exit decisions are observable to all players remaining in the game, these decisions may convey information on the aggregate state. Since types are correlated with state, this information is also relevant for learning each player's own type. We let $\xi_i^\theta(h^t)$ denote the probability with which player i exits after history h^t in state θ . Denoting by $\pi_i^\theta(h^t)$ the posterior probability with which player i is uninformed given history h^t , we have

$$\xi_i^\theta(h^t) := \sigma_i(h^t) \pi_i^\theta(h^t). \quad (1)$$

As long as $\xi_i^H(h^t) \neq \xi_i^L(h^t)$, other players learn about the true aggregate state by observing the exit decisions of player i . In contrast to the individual learning, the information content in observational learning from others is endogenous. First of all, if $\sigma_i(h^t) = 0$, then $\xi_i^H(h^t) = \xi_i^L(h^t) = 0$ and exit decisions are uninformative. Second, the amount of information depends on $\pi_i^\theta(h^t)$. It is important to note that $\pi_i^\theta(h^t)$ depends on $\sigma_i(h^s)$ for $s < t$ and hence is also endogenously determined. Subsection 3.2 describes how beliefs are updated based on observational learning from others. Subsection 3.3 combines the two forms of learning to derive the beliefs of an uninformed player in our model.⁸

⁸The reader more interested in the equilibrium analysis may jump to Section 4 and return to the present section as needed.

3.1 Private Learning

We start with the analysis of 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, i.e.

$$p_t := \Pr\{\text{"type of player } i \text{ is good"} \mid \text{"}i \text{ is uninformed in period } t\}.$$

If the player continues for another period, and receives a bad signal $\zeta = 0$, the new posterior p_{t+1} is obtained by Bayes' rule:

$$p_{t+1} = \frac{p_t(1 - \lambda\Delta)}{p_t(1 - \lambda\Delta) + 1 - p_t}. \quad (2)$$

The updating formula is essentially the same if the player knows the true aggregate state. We let p_t^θ denote the player's belief on her own type conditional on state θ . Using $p_0^\theta = \rho^\theta$, equation (2) gives us the formula for p_t^θ :

$$p_t^\theta = \frac{\rho^\theta(1 - \lambda\Delta)^t}{\rho^\theta(1 - \lambda\Delta)^t + (1 - \rho^\theta)}. \quad (3)$$

Notice that p_t^θ is a monotonic function of t and since it conditions on the state of the world, it will not be affected by learning from others.

To make the connection to the case with learning from others, we also describe the isolated player's beliefs on the aggregate state. We let q_t denote the probability that the individual player assigns on the state being H . By the law of iterated expectation,

$$p_t = q_t p_t^H + (1 - q_t) p_t^L. \quad (4)$$

Therefore

$$q_t = \frac{p_t - p_t^L}{p_t^H - p_t^L}, \quad (5)$$

and as a consequence, we see that the beliefs p_t and q_t move tightly together.

3.2 Observational Learning

To describe observational learning in our model, we consider for the moment how player i learns from the behavior of players $j \neq i$. We denote by $\widehat{q}_i(h^t)$ the belief of i on the aggregate state, when learning is based only on the behavior of the other players. Alternatively, we may think of $\widehat{q}_i(h^t)$ as the belief of player i as an *outside observer* to the game.

Recall that $\xi_j^\theta(h^t)$ denotes the probability with which an active player $j \in \mathcal{N}(h^t)$ exits at history h^t . If $\xi_j^L(h^t) > \xi_j^H(h^t) > 0$ and j does not exit, then i believes that j is more

likely to be informed, and that state H is relatively more likely. To describe the belief updating, we denote by $A_{-i}(h^t)$ the random vector containing the actions of all active players, excluding i , at history h^t . The probability of a given exit vector a_{-i}^t is then:

$$P_\theta(A_{-i}(h^t) = a_{-i}^t) = \prod_{j \neq i} ((1 - a_j^t) \xi_j^\theta(h^t) + a_j^t (1 - \xi_j^\theta(h^t))), \quad (6)$$

where we use shorthand notation P_θ to denote probability conditional on state:

$$P_H(\cdot) := \Pr(\cdot | \theta = H), \quad P_L(\cdot) := \Pr(\cdot | \theta = L).$$

After observing the exit vector a_{-i}^t , player i updates her belief $\hat{q}_i(h^t)$ according to Bayes' rule as follows:

$$\hat{q}_i(h^{t+1}) = \frac{\hat{q}_i(h^t) P_H(A_{-i}(h^t) = a_{-i}^t)}{\hat{q}_i(h^t) P_H(A_{-i}(h^t) = a_{-i}^t) + (1 - \hat{q}_i(h^t)) P_L(A_{-i}(h^t) = a_{-i}^t)}. \quad (7)$$

Note that $\hat{q}_i(h^{t+1})$ depends on $\xi_j^\theta(h^t)$ through (6), which in turn depends on $\pi_j^\theta(h^t)$ through (1). Therefore, in order to complete the description of observational learning in our model, we must also specify the evolution of $\pi_i^\theta(h^t)$, $i = 1, \dots, N$, for the fixed strategy profile σ . These posteriors change for two reasons within each period. First, at the beginning of each period exit decisions are realized. If player i continues, then the other players believe that i is more likely to be informed and update their beliefs using Bayes' rule as follows:

$$\pi_i^\theta(h^t) = \frac{\pi_i^\theta(h^t) (1 - \sigma_i(h^t))}{1 - \sigma_i(h^t) \pi_i^\theta(h^t)}, \quad \theta \in \{H, L\}. \quad (8)$$

Second, uninformed players become informed within the current period with probability $1 - p_t^\theta \lambda \Delta$ (conditional on not exiting). Combining these two steps, the updated belief after history h^{t+1} is:

$$\pi_i^\theta(h^{t+1}) = \frac{\pi_i^\theta(h^t) (1 - \sigma_i(h^t))}{1 - \sigma_i(h^t) \pi_i^\theta(h^t)} \cdot (1 - p_t^\theta \lambda \Delta), \quad \theta \in \{H, L\}. \quad (9)$$

3.3 Combined Learning

The remaining task is to combine the two forms of learning to derive the beliefs of the uninformed players in the game. The easiest way of doing this is to connect the information contained in $\hat{q}_i(h^{t+1})$ with the fact that i is uninformed.

Let π_t^θ denote the ex-ante probability with which a player (that stays in the game with probability 1) is uninformed in period t , conditional on state θ . The player is a bad type with probability $(1 - \rho^\theta)$, and all bad types remain uninformed. The player is a good

type with probability ρ^θ , and good types remain uninformed with probability $(1 - \lambda\Delta)$ in each period. Hence we have:

$$\pi_t^\theta = (1 - \rho^\theta) + \rho^\theta (1 - \lambda\Delta)^t. \quad (10)$$

We denote by $q_i(h^t)$ the belief of i on the aggregate state (conditional on being uninformed). This belief differs from $\widehat{q}_i(h^t)$ only to the extent that the private history of i affects her belief, and therefore the relationship between the two is given by Bayes' rule as follows:

$$q_i(h^t) = \frac{\widehat{q}_i(h^t) \pi_t^H}{\widehat{q}_i(h^t) \pi_t^H + (1 - \widehat{q}_i(h^t)) \pi_t^L}. \quad (11)$$

The belief of i on her own type follows now from the law of iterated expectation:

$$\begin{aligned} p_i(h^t) &= q_i(h^t) p_t^H + (1 - q_i(h^t)) p_t^L \\ &= \frac{\widehat{q}_i(h^t) \pi_t^H p_t^H + (1 - \widehat{q}_i(h^t)) \pi_t^L p_t^L}{\widehat{q}_i(h^t) \pi_t^H + (1 - \widehat{q}_i(h^t)) \pi_t^L}. \end{aligned} \quad (12)$$

where p_t^L and p_t^H are given by (3).

We end this section with two propositions that characterize learning in our model. First, we show that the likelihood of exit in state L is strictly larger than the likelihood of exit in H . Furthermore, the likelihood ratio of exit across the states is increasing over time. This guarantees that an exit is always informative about the aggregate state.

Equation (10) implies that the ex-ante likelihood ratio across states of being uninformed changes monotonically over time:

$$\frac{\pi_t^L}{\pi_t^H} > \frac{\pi_{t-1}^L}{\pi_{t-1}^H} > \dots > \frac{\pi_1^L}{\pi_1^H} = \frac{1 - \rho^L \lambda \Delta}{1 - \rho^H \lambda \Delta} > 1. \quad (13)$$

With this observation at hand, we can prove our first result on the informativeness of exits.

Proposition 1 *For any strategy profile σ , we have*

$$\frac{\xi_i^L(h^t)}{\xi_i^H(h^t)} \geq \frac{\pi_t^L}{\pi_t^H} > 1$$

for all $t > 0$ and h^t such that $\sigma_i(h^t) > 0$.

The second proposition ranks strategy profiles according to their informativeness. For a profile σ and a history h^t , we let the random variable $P_i^{t+1}(h^t, \sigma(h^t))$ denote the posterior of player i on her own type at the beginning of period $t + 1$, assuming that she

is uninformed at the beginning of period t . The randomness in the posterior arises from i 's private signal realization and the realized exit decisions of the players other than i . The following Proposition shows that higher exit probabilities by other players induce a mean preserving spread (in the sense of Rothschild & Stiglitz (1970)) on the posterior.

Proposition 2 *Take an arbitrary history h^t with $t > 0$ and two strategy profiles σ and σ' with $\sigma(h^s) = \sigma'(h^s)$ for $s = 0, \dots, t-1$. Then $P_i^{t+1}(h^t, \sigma(h^t))$ dominates $P_i^{t+1}(h^t, \sigma'(h^t))$ in the sense of second order stochastic dominance if $\sigma'_j(h^t) \geq \sigma_j(h^t)$ for all $j \neq i$ and $\sigma'_j(h^t) > \sigma_j(h^t)$ for some $j \neq i$.*

The economic content of this Proposition is rather immediate. Since all players' types are correlated with the state of the world θ , having maximal information on the state is also maximal information on an individual type. The total amount of information available to the players is captured by the vector of information types for the players, i.e. an enumeration of all players that are informed. A pure strategy profile $\sigma(h^t) = 1$ transmits all this information, since under this strategy players exit if and only if they are uninformed. The profile with $\sigma(h^t) = 0$ conveys no information. Any intermediate exit probability can be seen as a convex combination of these two signal structures, and it is to be expected that the combination with a higher weight on the informative signal is more informative with respect to the true informational state of a player.

4 Equilibrium Analysis

4.1 Isolated Player

It is again useful to start with the case of an isolated player. The decision problem of the isolated player is to choose whether to continue or exit at period t . Standard arguments show that the problem is a Markovian optimal stopping problem with the posterior probability $p = p_t$ as the state variable. We let $V_m(p)$ denote the value function of the isolated player. Stopping at posterior p yields a payoff of 0. If the player continues, she pays the cost $c\Delta$, and gets a good signal $\zeta = 1$ with probability $p\lambda\Delta$. In this case, the player learns that her expected payoff per period is $(\lambda v - c) \cdot \Delta$, and thus the value function jumps to

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

If the signal is $\zeta = 0$, then p falls to p_{t+1} according to (2). The Bellman equation for the optimal stopping problem can thus be written as:

$$V_m(p) = \max \left\{ 0, -c\Delta + p\lambda\Delta (v + \delta V^+(\Delta)) + (1 - p\lambda\Delta) \delta V_m \left(\frac{p(1 - \lambda\Delta)}{p(1 - \lambda\Delta) + (1 - p)} \right) \right\}. \quad (14)$$

Since the expected payoff is increasing in p , the optimal policy is to stop as soon as p falls below a threshold level, which we denote by $p^*(\Delta)$. Standard arguments establish that the value function $V_m(p)$ is increasing, convex and continuous in p . The threshold $p^*(\Delta)$ is obtained from (14) by setting $V_m(p^*(\Delta)) = 0$:

$$p^*(\Delta) = \frac{c}{\lambda(v + \delta V^+(\Delta))}. \quad (15)$$

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 all signals so far have been bad:

$$t^*(\Delta) := \min \{t \in \mathbb{N} \mid p_t \leq p^*(\Delta)\}.$$

We denote the optimal strategy of the isolated player by

$$a_m(p_t) = \begin{cases} 1 & \text{if } p_t > p^*(\Delta), \\ 0 & \text{if } p_t \leq p^*(\Delta). \end{cases}$$

4.2 Symmetric Equilibrium

In this subsection, we show that the exit game with observational learning has a unique symmetric equilibrium.⁹ Furthermore, the equilibrium value functions of the individual players can be written as functions of their belief on their own type only. With symmetric strategies all uninformed players have identical beliefs, and therefore we drop the subscripts i from the beliefs and the strategies of the uninformed players. In particular, we let $\sigma(h^t)$ denote the probability with which each uninformed player exits at history h^t in symmetric equilibrium.

We start by showing that if a symmetric equilibrium exists for the stopping game, then the equilibrium value function $V(h^t)$ is closely related to the value function of the isolated player. The key observation for this result is that as long as $\sigma(h^t) = 0$, there is no observational learning and thus the information available in the game is identical to the information available to the isolated player. On the other hand, when $\sigma(h^t) > 0$,

⁹The game has also asymmetric equilibria, where the players act in a predetermined order conditioning their actions on the outcomes of the previous moves by the other players. Since the properties of such equilibria are essentially similar to the herding models with exogenous order of moves, we do not discuss them further (details about asymmetric equilibria are available from authors upon request).

the players can learn from each other, but their value must be zero since they exit with a positive probability.¹⁰

Lemma 1 *For any symmetric equilibrium of the exit game,*

$$V(h^t) = V_m(p(h^t)).$$

With the help of Lemma 1 we can derive a symmetric equilibrium strategy profile recursively. To see this, note that if a symmetric equilibrium is given for periods $0, \dots, t-1$, we can calculate the beliefs of uninformed players at history h^t as explained in Section 3. Consider then exit probabilities at history h^t . By Lemma 1, the payoff for the next period is given by $V(h^{t+1}) = V_m(p(h^{t+1}))$, and therefore, all we have to do is to find an exit probability $\sigma(h^t)$ that induces a probability distribution for $p(h^{t+1})$ that makes the players indifferent between exiting and staying. This indifference condition must equate the discounted expected value for the next period with the cost of staying for one period, so we can write it as:

$$\delta \mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t))) = c(h^t) \cdot \Delta,$$

where we use notation $c(h^t)$ to denote cost of staying net of expected payoff per time unit:

$$c(h^t) := c - p(h^t) \lambda v.$$

The next Lemma shows that increasing the exit probabilities for the current period increases the players' incentive to stay. This result follows from two observations. First, by Proposition 2, increasing the exit probability for the current period induces a mean preserving spread for the next period belief $P^{t+1}(h^t, \sigma(h^t))$. Second, we know from Subsection 4.1 that the isolated player's value function V_m is convex. The lemma also guarantees that this monotonicity property is strict at the point of indifference, which is essential for the uniqueness of symmetric equilibrium.

Lemma 2 *The expected continuation payoff $\mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t)))$ is weakly increasing in $\sigma(h^t)$. Furthermore, for each h^t such that $p(h^t) < p^*(\Delta)$, there is at most one exit probability $\sigma(h^t)$ satisfying*

$$\delta \mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t))) = c(h^t) \cdot \Delta.$$

¹⁰In Section 7, we discuss how this result is modified in extensions of the current model where the active players may have heterogenous private histories.

Lemma 2 guarantees that for each h^t (except for the knife-edge case $p(h^t) = p^*(\Delta)$), at most one exit probability can make the players indifferent between exiting and staying.¹¹ However, if the players are so pessimistic that even the information transmitted by the pure strategy profile $\sigma(h^t) = (1, \dots, 1)$ is not sufficient to compensate for the loss $c(h^t)\Delta$ of waiting for one more period, then it is a dominant action for all the uninformed players to exit with probability 1. When this happens, we say that the game *collapses*.

With these preliminaries, we are ready to prove the existence and the uniqueness of a symmetric equilibrium.¹²

Theorem 1 *The stopping game has a unique symmetric equilibrium where the exit probability at history h^t is given by:*

$$\sigma(h^t) = \begin{cases} 0 & \text{if } \delta\mathbb{E}V_m(P^{t+1}(h^t; 0)) > c(h^t)\Delta, \\ \sigma^*(h^t) \in [0, 1] & \text{if } \delta\mathbb{E}V_m(P^{t+1}(h^t; 0)) \leq c(h^t)\Delta \leq \delta\mathbb{E}V_m(P^{t+1}(h^t; 1)), \\ 1 & \text{if } \delta\mathbb{E}V_m(P^{t+1}(h^t; 1)) < c(h^t)\Delta, \end{cases}$$

where $\sigma^*(h^t)$ solves

$$\delta\mathbb{E}V_m(P^{t+1}(h^t; \sigma^*(h^t))) = c(h^t)\Delta.$$

The symmetric equilibrium has a simple structure. Whenever the players' beliefs on their own type are above the threshold of the isolated player, i.e. $p(h^t) > p^*(\Delta)$, then $\delta\mathbb{E}V_m(P^{t+1}(h^t; 0)) > c(h^t)\Delta$ and thus the equilibrium actions coincide with those prescribed by the optimal decision rule of the isolated player (i.e. stay). On the other hand, if the players are very pessimistic, then $\delta\mathbb{E}V_m(P^{t+1}(h^t; 1)) < c(h^t)\Delta$, and again equilibrium actions coincide with the isolated player (i.e. exit with probability one). With intermediate beliefs equilibrium behavior differs from isolated player: in equilibrium the exits take place with a probability that exactly balances the players' incentives to exit and wait, whereas an isolated player would exit with probability one.

We will obtain a sharper characterization of the symmetric equilibrium when we decrease the time lag between successive periods towards zero. This will be done in Section 6.

5 Information Aggregation in Large Games

In this section, we analyze information aggregation in the symmetric equilibrium of the game as the number of players grows without bound. For this purpose, the natural point

¹¹For the non-generic case where $p(h^t) = p^*(\Delta)$ after some history h^t , all randomizations $\sigma(h^t)$ that result in posteriors $p(h^{t+1}) \leq p^*(\Delta)$ for all a^t are compatible with equilibrium.

¹²The uniqueness is modulo the multiplicity at $p(h^t) = p^*$ as explained in the previous footnote.

of comparison is the case where the players share all information with each other. If the number of players is large and all information is pooled, then the (weak) law of large numbers implies that the players can determine the true aggregate state with arbitrarily high accuracy. Nevertheless, idiosyncratic uncertainty about player types remains: conditional on the state, each player is still uncertain about her own type. As a result, the efficient benchmark in terms of information aggregation for large games is the one where all the uninformed players know the aggregate state θ . In state θ , an uninformed player believes that she is a good type with probability p_t^θ . Therefore it is optimal for her to exit as soon as p_t^θ falls below $p^*(\Delta)$. We denote the efficient exit period in state θ by $t_\theta^*(\Delta)$:

$$t_\theta^*(\Delta) := \min \{t : p_t^\theta \leq p^*(\Delta)\}, \theta = H, L.$$

Clearly, we have $t_L^*(\Delta) < t^*(\Delta) < t_H^*(\Delta)$.¹³

The main result of this section is Theorem 2, which says that by decreasing the period length, we eliminate the possibility that a large number of players exit too early relative to this efficient benchmark. This means that for large games, information is aggregated efficiently in state $\theta = H$ because $t_H^*(\Delta)$ is an upper bound for all exit times of the uninformed players. However, if $\theta = L$, information aggregation fails: all players exit too late in expectation.

Since we vary the number of players N and the period length Δ while keeping all the other parameters of the model fixed, we denote by $\Gamma(\Delta, N)$ the game parametrized by Δ and N . We denote by $X(h^t)$ the number of players that exit the game at history h^t in the unique symmetric equilibrium of the game:

$$X(h^t) := n(h^t) - n(h^{t+1}).$$

As a first step towards Theorem 2, we consider the effect of a large number of exits on the beliefs. We already showed in Proposition 1 that individual exit probabilities are different across the two states, which allows the players to make inferences based on observed exits. It is therefore natural to expect that if a large number of players exit, then all the remaining players should have accurate beliefs on the aggregate state. Proposition 3 shows that this must indeed be the case with a high probability. The idea in the proof is to follow the belief of an outside observer to the game, and to show that this belief must converge to truth at a high rate as the number of exits grows. The argument uses techniques developed in Fudenberg & Levine (1992).

¹³Unless ρ^L and ρ^H are very close to each other and Δ is large, in which case we may have $t_L^*(\Delta) = t_H^*(\Delta)$. In that case the optimal action does not depend on state, and observational learning is irrelevant.

Proposition 3 For all $\varepsilon > 0$, there is some $K \in \mathbb{N}$ such that

$$P_H\{h^\infty : n(h^t) \leq N - K \text{ and } q(h^t) < 1 - \varepsilon \text{ for some } h^t \in h^\infty\} < \varepsilon, \quad (16)$$

$$P_L\{h^\infty : n(h^t) \leq N - K \text{ and } q(h^t) > \varepsilon \text{ for some } h^t \in h^\infty\} < \varepsilon, \quad (17)$$

for any game $\Gamma(\Delta, N)$.

A couple of remarks are in order. First, Proposition 3 holds for all strategy profiles σ (equilibrium or not) as long as some private information has been accumulated before first exits. Second, the bound K for the number of exits in the Proposition is independent of Δ and N . By increasing N , we can make sure that the state is revealed if an arbitrarily small fraction of players exit.

Proposition 3 implies that once K players have exited and $\theta = H$, then with a high probability, no further exits take place before $t_H^*(\Delta)$. This would suggest that the total number of suboptimally early exits must be bounded. However, we must also consider the possibility that an unbounded number of players exit within *a single period* before they learn the true state. Our second step towards Theorem 2 is to show that by reducing period length Δ , we can eliminate this possibility. This is established in Proposition 4 below.

We need some notation to keep track of the passage of real time as we vary Δ .¹⁴ Let τ_θ denote the efficient exit time corresponding to state θ in the limit $\Delta \rightarrow 0$:

$$\tau_\theta := \lim_{\Delta \rightarrow 0} [t_\theta^*(\Delta) \cdot \Delta], \theta = H, L.$$

To link real time to the corresponding period of a discrete time model, we define $t(\tau, \Delta)$ as the last period before an arbitrary real time τ :

$$t(\tau, \Delta) := \max\{t : t \cdot \Delta \leq \tau\}. \quad (18)$$

Proposition 4 For all $\tau < \tau_H$ and $\varepsilon > 0$, there are constants $\bar{\Delta} \in \mathbb{R}^+$ and $K \in \mathbb{N}$ such that

$$P_H\{h^\infty : X(h^t) > K \text{ for some } t \leq t(\tau, \Delta)\} < \varepsilon,$$

for any game $\Gamma(\Delta, N)$ with $\Delta < \bar{\Delta}$.

The proof of Proposition 4 is lengthy, but the intuition is straight-forward. If the players were to adopt a strategy that induces a large number of exits with non-negligible

¹⁴This notation will also be useful in the following Section where the continuous time limit of the model is considered.

probability within a single period, then this would generate a very informative signal about the state. For all $\tau < \tau_H$, the value of such a signal is positive. If the waiting cost is low enough, then all the players would have a strict preference to observe the signal rather than exit contradicting the hypothesized positive probability of exits.

Combining Propositions 3 and 4 gives us the theorem that bounds the total number of suboptimally early exits in the game. This result implies that in the double limit where we increase the number of players and reduce the period length, the fraction of players that exit suboptimally early shrinks to zero.

Theorem 2 *For all $\tau < \tau_H$ and $\varepsilon > 0$, there are constants $\bar{\Delta} \in \mathbb{R}^+$ and $K \in \mathbb{N}$ such that*

$$P_H \left\{ h^\infty : \sum_{t=0}^{t(\tau, \Delta)} X(h^t) > K \right\} < \varepsilon,$$

for any game $\Gamma(\Delta, N)$ with $\Delta < \bar{\Delta}$.

This theorem is important for us in two respects. First, as we will see in the next section, it allows us to compute explicitly the statistical properties of the equilibrium path in the limit where Δ is small and N is large. Second, it carries a central message about information aggregation in our model: information cannot be fully aggregated before the optimal stopping time in the true state.

Perhaps the most illuminating special case is the case of perfectly correlated private types (i.e. $\rho^H = 1$ and $\rho^L = 0$) where full information on the aggregate state is sufficient to determine all players' payoff types. In this case $\tau_H = \infty$, and Theorem 2 implies that in state H almost all players remain in the game forever. In state L , all players exit the game eventually, but they do so later than they would in the absence of observational learning.

Although we have assumed symmetry throughout this section, there is very little that depends on this restriction. The proof of Proposition 3 is valid for any asymmetric equilibrium strategy profile as well. The proof of Proposition 4 uses symmetry in two lemmas (Lemma 6 and Lemma 8 in Appendix B). However, even there symmetry is used for notational convenience (the number of exits within a period is binomially distributed, which leads more easily to the desired results).

6 Exit Waves

In this section, we characterize the symmetric equilibrium in the limit as $\Delta \downarrow 0$. We have several reasons for this. The first reason is substantive. In a model with endogenous

timing decisions, it is important to know if the results depend on an exogenously imposed reaction lag Δ . Second, it turns out that the inherent dynamics of the model are best displayed in the limit: information aggregation happens in randomly occurring bursts of sudden activity. We call these bursts of activity *exit waves*. Third, when we also let $N \rightarrow \infty$, we can compute the statistical properties of the equilibrium path in an explicit form.

We may view the public history h^∞ generated by the symmetric equilibrium $\sigma(\Delta, N)$ in the game $\Gamma(\Delta, N)$ with period length Δ and N players from a slightly different angle. Suppose that the players are to be treated anonymously. Then the vector $\mathbf{t}(\Delta, N) = (t_1(\Delta, N), \dots, t_N(\Delta, N))$ where $t_k(\Delta, N)$ indicates the period in which k^{th} exit took place gives a full account h^∞ . The profile $\sigma(\Delta, N)$ induces a probability distribution on \mathbb{R}^N on instants of exit measured in continuous time τ (that is, k^{th} exit takes place at time $\tau_k = t_k(\Delta, N) \cdot \Delta$). We denote this distribution, conditional on state θ , by $F_{\Delta, N}^\theta(\tau)$. We investigate the limiting distribution

$$F_N^\theta(\tau) = \lim_{\Delta \downarrow 0} F_{\Delta, N}^\theta(\tau),$$

where the convergence is taken to be in the sense of weak convergence. Observational learning then results from the differences between $F_N^H(\tau)$ and $F_N^L(\tau)$.

In Subsection 6.1 we keep the number of players fixed at N . We show that when there was no exit in the previous period, the probability of an exit within the current period is of the order Δ . This means that exits arrive according to a well defined hazard rate, and we say that the game is in the *flow mode*. On the other hand, if there was an exit in the previous period, then the probability of an exit in the current period is bounded away from zero, and we say that the game is in an *exit wave*.

In Subsection 6.2, we consider the limiting distributions

$$F^\theta(\tau) = \lim_{N \rightarrow \infty} F_N^\theta(\tau)$$

defined on the set of sequences of exit times $\{\tau_k\}_{k=1}^\infty$. In particular, we compute the distributions for the first K exit instants and we also calculate the probability of the event that the market collapses by time instant τ , i.e. the probability of the event $\{\tau_k \leq \tau \text{ for all } k\}$. We make use of Poisson approximations and Theorem 2 when computing the size of the exit events and the probability of a collapse given that an exit event started.

6.1 The Structure of the Symmetric Equilibrium

In this subsection, we keep the number of players N fixed. Since we are interested in the limit $\Delta \rightarrow 0$, we parametrize the game and its histories with the period length Δ . We

say that the game is in the *flow mode* at history h^t if no players exited at history h^{t-1} , i.e. if $X(h^{t-1}) = 0$. The game is in an *exit wave* at history h^t if $X(h^{t-1}) > 0$. Finally, we say that the game *collapses* at history h^t if $\sigma(h^t) = 1$. Collapse is an absorbing state: since all uninformed players exit, the game is effectively over, and $\pi(h^s) = 0$ for all $s > t$. This means that for a game with a given Δ , we have three mutually exclusive sets of histories, corresponding to flow mode, exit wave, and collapse, respectively:

$$\begin{aligned} H^f(\Delta) & : = \{h^t : X(h^{t-1}) = 0 \text{ and } \pi(h^t) > 0\} \\ H^w(\Delta) & : = \{h^t : X(h^{t-1}) > 0 \text{ and } \pi(h^t) > 0\} \\ H^c(\Delta) & : = \{h^t : \pi(h^t) = 0\} \end{aligned}$$

In order to relate the discrete decision periods to real time instants, we define

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

where $p^*(\Delta)$ and $t^*(\Delta)$ denote the belief threshold and the corresponding exit time as defined in Section 4.1.

We start by showing that the beliefs of the uninformed players are qualitatively different in the two active modes. When the game is in the flow mode and Δ is small, beliefs are close to p^* while in an exit wave, they are bounded away from p^* .

Lemma 3 *i) For all $\varepsilon > 0$, there is a $\bar{\Delta} > 0$ such that*

$$p(h^t) \in (p^* - \varepsilon, p^* + \varepsilon)$$

for all $h^t \in H^f(\Delta)$, $t \geq t^(\Delta)$, $\Delta < \bar{\Delta}$.*

ii) There is an $v > 0$ and $\bar{\Delta} > 0$ such that

$$p(h^t) < p^* - v$$

for all $h^t \in H^w(\Delta)$, $\Delta < \bar{\Delta}$.

The following Proposition shows that active players also behave differently in the two modes.

Proposition 5 *i) There is a $\kappa > 0$ such that*

$$P_H(X(h^t) > 0) < P_L(X(h^t) > 0) < \kappa\Delta \text{ for all } h^t \in H^f(\Delta), \Delta > 0.$$

ii) There is a $\underline{p} > 0$ and a $\bar{\Delta} > 0$ such that

$$P_L(X(h^t) > 0) > P_H(X(h^t) > 0) > \underline{p} \text{ for all } h^t \in H^w(\Delta), 0 < \Delta < \bar{\Delta}.$$

The first claim in the above Proposition justifies our use of the term *flow mode*. The flow mode comes to an end with a well-defined hazard rate. The actual computation of the equilibrium hazard rate is not hard in principle. Nevertheless, the formula will depend on the evolution of $\pi^\theta(h^t)$ and it is not possible to give a closed-form solution for the continuous-time limit of the updating formula (9).¹⁵ In the following subsection, we compute the explicit hazard rate in the limit as $N \rightarrow \infty$.

The second point to note is that for all finite N , any exit wave ends in a finite number of periods. Hence the real time that an exit wave lasts is bounded from above by ΔN and vanishes as $\Delta \rightarrow 0$. As a result, we may view the limit exit waves as (rather complicated) randomization events between the flow mode and collapse.

Finally, every exit wave results in a collapse with a strictly positive probability. Since an exit wave takes only a vanishing amount of real time, learning from own experience during the wave can be ignored. Recall that $p(h^t)$ is a martingale, and its value in exit wave mode is bounded away from p^* by part ii) of Lemma 3. Therefore, we conclude that a return to the flow mode cannot happen with probability 1. If the number of players is small, then the first exit starting a wave may in fact lead to an immediate collapse. Then the exit wave lasts only one period and ends up in collapse with probability one. If this is not the case, then a martingale argument establishes that the game must return to the flow mode with a strictly positive probability.

6.2 Exit Events in Large Games

The large game limit $N \rightarrow \infty$ simplifies the computations for a number of reasons. First, we can use Poisson approximations of the Binomial distribution for the number of exits within each period of an exit wave. Second, conditional on the game not having collapsed, we know that $\pi_i^\theta(h^t) \approx \pi_t^\theta$ for all t , which allows us to use the continuous time limit of equation (10) instead of equation (9) to compute the conditional probabilities for the players to be uninformed. This follows from the fact that with a large number of players, each individual player exits with a negligible probability (until collapse). Third, we can apply Theorem 2, which implies that the game collapses in state $\theta = H$ before τ_H with a vanishing probability as $N \rightarrow \infty$.

¹⁵The computation is somewhat complicated because the updating of $\pi^\theta(h^t)$ depends on the equilibrium randomization probabilities $\sigma(h^t)$.

Let $p^\theta(\tau)$ and $\pi^\theta(\tau)$ denote the continuous time limits of (3) and (10):

$$\begin{aligned} p^\theta(\tau) &: = \frac{\rho^\theta e^{-\lambda\tau}}{(1 - \rho^\theta) + \rho^\theta e^{-\lambda\tau}}, \\ \pi^\theta(\tau) &: = (1 - \rho^\theta) + \rho^\theta e^{-\lambda\tau}, \end{aligned}$$

and let $q^*(\tau)$ denote the belief on aggregate state that corresponds to the critical private belief p^* :

$$q^*(\tau) := \frac{p^* - p^L(\tau)}{p^H(\tau) - p^L(\tau)}, \quad \tau_L \leq \tau \leq \tau_H.$$

Note that $q^*(\tau)$ is strictly increasing within $[\tau_L, \tau_H]$ and $q^*(\tau_L) = 0$ and $q^*(\tau_H) = 1$.

We compute first the hazard rate of exits in the flow mode. In particular, assume that k players have exited the game at real times τ_1, \dots, τ_k , and the game is in flow mode at real time τ . Using the fact that the likelihood ratio of exit across states is given by $\pi^L(\tau)/\pi^H(\tau)$, and the fact that the belief of an uninformed player must stay close to p^* as long as no player exits (as required by Lemma 3), we can calculate the hazard rate with which an additional player exits:

Proposition 6 *In the limit $N \rightarrow \infty$, the instantaneous hazard rate of $k + 1^{\text{st}}$ exit at some $\tau \in (\tau_k, \tau_H)$, conditional on the first k exit times τ_1, \dots, τ_k , is given by*

$$\frac{f_{k+1}^\theta(\tau | \tau_1, \dots, \tau_k)}{1 - F_{k+1}^\theta(\tau | \tau_1, \dots, \tau_k)} = \pi^\theta(\tau) \lambda \frac{p^*(1 - p^*)(p^H(\tau) - p^L(\tau))}{(p^* - p^L(\tau))(p^H(\tau) - p^*)(\pi^L(\tau) - \pi^H(\tau))} := \psi^\theta(\tau). \quad (19)$$

The notation $\psi^\theta(\tau)$ is appropriate, because the expression (19) does not depend on τ_1, \dots, τ_k . Note that (19) applies also in the case where no player has yet exited. Hence, irrespective of the number of players that have already exited, $\psi^\theta(\tau)$ gives the hazard rate with which an exit wave starts at time τ , conditional on the game being in flow mode.

Every exit wave leads either to collapse or a return to flow mode. With a large number of players, it is easy to compute the probabilities with which either possibility occurs. To see this, consider the posterior after the first exit of an exit wave that takes place at real time τ . Since the belief just before the exit is $q^*(\tau)$ and the exit reveals one player to be uninformed, the posterior after the first exit is given by Bayes' rule:

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

By Theorem 2, the game returns to the flow mode with a probability that converges to 1 as $N \rightarrow \infty$ in state $\theta = H$. Therefore, if the game collapses, $q(\tau+) = 0$. On the other hand, we know from Lemma 3 that if the game returns to the flow mode, we

have $q(\tau+) = q^*(\tau)$. Let $\phi^\theta(\tau)$ denote the probability of collapse given an exit event at $\tau < \tau_H$, conditional on state θ . By Theorem 2,

$$\phi^H(\tau) = 0,$$

so the probability of collapse estimated by a player with belief $q^-(\tau)$ is $(1 - q^-(\tau))\phi^L(\tau)$. Therefore, by the martingale property of belief for this player we have:¹⁶

$$q^-(\tau) = (1 - (1 - q^-(\tau))\phi^L(\tau))q^*(\tau),$$

which gives

$$\phi^L(\tau) = \frac{q^*(\tau) - q^-(\tau)}{q^*(\tau)(1 - q^-(\tau))} = \frac{\pi^L(\tau) - \pi^H(\tau)}{\pi^L(\tau)}. \quad (20)$$

Since (19) gives the hazard rate with which an exit wave starts, and (20) gives the probability with which a given exit wave leads to collapse, we get the hazard rate of collapse by multiplying them:

Corollary 1 *In the limit $N \rightarrow \infty$, the instantaneous hazard rate of collapse at time $\tau \in (\tau_L, \tau_H)$, conditional on state, and conditional on being in flow mode at τ , is:*

$$\chi^H(\tau) = 0, \quad (21)$$

$$\chi^L(\tau) = \lambda \frac{p^*(1 - p^*)(p^H(\tau) - p^L(\tau))}{(p^* - p^L(\tau))(p^H(\tau) - p^*)}. \quad (22)$$

Note that equations (21) and (22) tell us how an outside observer learns from the actions of the players. Since the game collapses only in state L , an outside observer becomes gradually more optimistic over time about the aggregate state if there is no collapse. By Bayes' rule, the belief dynamics of the outside observer is given by:

$$\frac{d\hat{q}(\tau)}{d\tau} = \hat{q}(\tau)(1 - \hat{q}(\tau))\chi^L(\tau).$$

If $\rho^H < 1$, $p^H(\tau)$ falls to p^* at $\tau = \tau_H$. We see from (22) that

$$\lim_{\tau \rightarrow \tau_H} \chi^L(\tau) = \infty.$$

In words, the hazard rate of collapse in state L explodes at τ_H . This is to be expected, because τ_H is an upper bound for all exit times. As a result, the state is eventually fully revealed to the outside observer: if the game collapses strictly before τ_H , then $\theta = L$, otherwise $\theta = H$.

¹⁶If the game has $N < \infty$ players, then collapse will take place at a posterior $q^C > 0$ and as a consequence, the probability of a collapse is higher than in equation (20).

Corollary 1 applies also to the special case of perfectly correlated types (i.e. $\rho^H = 1$ and $\rho^L = 0$). In that case, $p^H(\tau) \equiv 1$ and $p^L(\tau) \equiv 0$, and hence Equation (22) reduces to $\chi^L(\tau) = \lambda$. To understand this simple formula, note that an uninformed player learns fully her own type from either receiving a good private signal (in which case she must be a good type) or by observing the game collapse (in which case she must be a bad type). Lemma 3 requires that her belief stay constant as long as none of these two events occur. This is possible only if both events arrive at the same rate conditional on state. Note that in this special case $\tau_H = \infty$, and hence if $\theta = H$, the game never collapses and the outside observer learns only asymptotically the true state.

We end this section by describing the sequence of events *within* a given exit wave that takes place at real time τ . We use index $s = 1, 2, \dots$ to refer to the consecutive periods within the exit wave and S to denote the total number of periods within the wave. Let q_s denote the belief in the s^{th} period of the wave, and let X_s denote the number of exits at that period. Note that since we are considering the limit $\Delta \rightarrow 0$, the duration of the exit wave in real time is zero in the limit.

Fix a period s and the corresponding belief q_s . Lemma 3 implies that we must have $q_s < q^*(\tau)$. On the other hand, the same lemma implies that if s is the last period of the exit wave (that is, no player exits), then we must have $q_{s+1} = q^*(\tau)$.

Proposition 7 *Consider period s of an exit wave taking place at time τ . As $N \rightarrow \infty$, X_s^θ converges in distribution to a Poisson random variable with parameter:*

$$\frac{\pi^\theta(\tau)}{(\pi^L(\tau) - \pi^H(\tau))} \log \left(\frac{q^*(\tau)}{(1 - q^*(\tau))} \cdot \frac{(1 - q_s)}{q_s} \right) \text{ for } \theta \in \{H, L\}. \quad (23)$$

If the realized number of exits is $X_s = k$, the next period belief is:

$$q_{s+1}(k) := \frac{(\pi^H(\tau))^k q^*(\tau)}{(\pi^H(\tau))^k q^*(\tau) + (\pi^L(\tau))^k (1 - q^*(\tau))}. \quad (24)$$

Notice from (24) that the number of exits in the previous stage is a sufficient statistic for the belief in the current stage, which in turn by (23) is a sufficient statistic for current stage randomization probabilities. Hence the limiting process of beliefs on the state of the world can be thought of as a non-homogenous Poisson random walk on the positive integers.

The exit event taking place at real time instant τ reverses to flow mode if there is some s such that $X_s = 0$. The Poisson approximation formula above in equation (23), gives the probability that the exit event ends after s periods. The total number of exits

at real time instant τ is given by

$$X(\tau) = 1 + \sum_{s=1}^S X_s.$$

7 Discussion

We analyzed a stopping game where the players base their decisions on their privately acquired information and on the behavior of the other players in a similar situation. The ability to observe the actions of others makes the players more likely to postpone their actions. But this, in turn, reduces the informativeness of their actions, thus reducing the incentives to wait. The symmetric equilibrium balances these effects and leads to aggregate delays and randomly arriving exit waves. We showed that even when the number of players gets large, aggregate uncertainty persists in equilibrium. Information is aggregated gradually until a sudden collapse leads to full revelation of the aggregate state.

We kept the model as simple as possible in order to highlight the interplay between individual and social learning. In the following subsections, we discuss a number of extensions of the model.

7.1 Heterogenous Private Values

In the main model, all players have an identical opportunity cost c and an identical value from the signal v . We could have equally well assumed that c_i is private information and drawn from a common distribution $F(c_i)$ at the beginning of the game. With this specification, the game has a pure Perfect Bayesian Equilibrium in symmetric strategies. An isolated player has an optimal policy characterized by a stopping threshold $p^*(\Delta; c_i)$ and an optimal value given by $V_m(p(h^t); c_i)$. With heterogenous private values, the conclusion of Lemma 1 holds only for players with high c_i . For each history (h^t, c_i) , there exists a cutoff cost type $c(h^t)$ such that the symmetric equilibrium value $V(h^t; c_i)$ is equal to the expected payoff of the isolated player for all $c_i > c(h^t)$.

Adjusting the cutoff level $c(h^t)$ plays the same role as changing the equilibrium randomization probabilities $\sigma(h^t)$ in the main model. It is no longer true that all players get the same payoff as they would as isolated players (the players with lower c benefit from observational learning). Nevertheless, the characterization of the equilibrium path remains almost unchanged.

It is perhaps worth noticing that with heterogenous costs, the isolated players' optimal stopping times are continuously distributed as a function of c_i . In the case where the

individual types are perfectly correlated (i.e. $\rho^H = 1$ and $\rho^L = 0$) and players observe each others' actions, almost all players exit the game at the same time if the number of players is large.

7.2 Bounded Signals

We stressed the importance of our assumption that the players do not become unboundedly pessimistic as a result of their private learning. Yet we allow for the fact that signals are perfectly informative for the good types. In this subsection, we argue that this assumption is purely for notational simplicity.

A more general model would allow for two Poisson rates of signal arrivals, λ^G if the player is a good type, and λ^B for a bad type. The players' private histories are now summed by k_i^t , the number of signals that player i has observed up to period t . The optimal stopping policy of the isolated player is still characterized by a cutoff rule $p^*(\Delta)$. Denote the symmetric equilibrium payoff in this case by $V(h^t, k_i^t)$. In analogy to Lemma 1, we can show that for all h^t such that $p(h^t, k_i^t) < p^*(\Delta)$ for some active players i , there is a cutoff level of private histories $k(h^t)$ such that all players with $k_i^t < k(h^t)$ earn the same expected continuation payoff as an isolated player with the same belief of her own type.

The equilibrium randomization probabilities are determined through the requirement that the players with exactly $k(h^t)$ signals must be indifferent between continuing and exiting. Those players with fewer than $k(h^t)$ signals exit with probability one if they have not exited already.

In the modified model, all players are ex ante identical, but they become heterogenous due to their different private experiences. Since only the most pessimistic players exit with positive probability, the ex ante equilibrium value in the model with observational learning exceeds the ex ante expected payoff of the isolated player. Furthermore, with perfectly correlated private types, almost all players exit the game at the same time when the number of players is large.

7.3 Payoffs Depending on Aggregate State

We could allow the value of the signal in aggregate state θ be v^θ . In this case, the value of being informed, V^θ , would also depend on the state.¹⁷ As long as the signals accumulate at the same rate across the two states, the analysis remains similar to the main model.

¹⁷The different values across states could reflect differences in demand conditions in monopolistically competitive markets.

Lemma 1 remains unchanged, and therefore the analysis is also almost identical to the current case.

The main difference compared to the current model is in the characterization of the indifference condition of the uninformed players. In the current paper, indifference requires (in the continuous time limit) that $p(\tau) = p^*$. In the extended model the analogous requirement would be that

$$\lambda[q(\tau)p^H(\tau)V^H + (1 - q(\tau))p^L(\tau)V^L] = c,$$

where

$$V^\theta = v^\theta + \frac{\lambda v^\theta - c}{r}.$$

7.4 Payoff Externalities

Finally, a more challenging extension would be to incorporate payoff externalities in the model. The payoff could for example depend on the number of players present in the market. It seems to us that beyond the two-player case, quite different analytical techniques would be needed to cover this case. Lemma 1 has no analogue in this extension, and as a result, the analysis will have to be quite different. In our view, this is an interesting and challenging direction for further research.

Appendix A: Shorter Proofs

Proof of Proposition 1. Since $\xi_i^\theta(h^t) = \sigma_i(h^t) \pi_i^\theta(h^t)$, all we have to do is to show that

$$\frac{\pi_i^L(h^t)}{\pi_i^H(h^t)} \geq \frac{\pi_t^L}{\pi_t^H} \quad (25)$$

for all $t > 0$ (note that $\pi_t^L/\pi_t^H > 1$ follows from (13)).

We use induction. As an induction hypothesis, assume that (25) holds for some $t \geq 0$. Using (9) and (25), we then have

$$\begin{aligned} \frac{\pi_i^L(h^{t+1})}{\pi_i^H(h^{t+1})} &= \left(\frac{\pi_i^L(h^t)}{\pi_i^H(h^t)} \right) \left(\frac{1 - \sigma_i(h^t) \pi_i^H(h^t)}{1 - \sigma_i(h^t) \pi_i^L(h^t)} \right) \left(\frac{1 - p_t^L \lambda \Delta}{1 - p_t^H \lambda \Delta} \right) \\ &\geq \frac{(1 - p_t^L \lambda \Delta) \pi_t^L}{(1 - p_t^H \lambda \Delta) \pi_t^H}. \end{aligned} \quad (26)$$

On the other hand, using (3) and (10), we have:

$$\pi_{t+1}^\theta = (1 - p_t^\theta \lambda \Delta) \pi_t^\theta, \theta = H, L.$$

Combining this with (26) gives us the induction step:

$$\frac{\pi_i^L(h^{t+1})}{\pi_i^H(h^{t+1})} \geq \frac{\pi_{t+1}^L}{\pi_{t+1}^H}.$$

Noting that $\pi_i^L(h^0) = \pi_i^H(h^0) = \pi_0^L = \pi_0^H = 1$ gives us:

$$\frac{\pi_i^L(h^0)}{\pi_i^H(h^0)} \geq \frac{\pi_0^L}{\pi_0^H},$$

and therefore, the proof by induction is complete. ■

Proof of Proposition 2. Construct an experiment X_i on $\Theta = \{H, L\}$ with outcomes in $S^{X_i} = \{0, 1\}$. The joint probabilities on the states and outcomes are given by the following stochastic matrix P^{X_i}

P^{X_i}	$\theta = H$	$\theta = L$
$s^{X_i} = 1$	$1 - \sigma_i(h^t) \pi_i^H(h^t)$	$1 - \sigma_i(h^t) \pi_i^L(h^t)$
$s^{X_i} = 0$	$\sigma_i(h^t) \pi_i^H(h^t)$	$\sigma_i(h^t) \pi_i^L(h^t)$

We interpret the event $\{\theta = H\}$ as the event that the state is good and the event $\{s^{X_i} = 1\}$ as the decision of player i to stay in the game. The joint probability over (θ, s^{X_i}) simply reflects the conditional exit probabilities given strategy σ .

Consider next another experiment Y_i on Θ with outcomes in $S^{Y_i} = \{0, 1\}$ and the associated stochastic matrix P^{Y_i}

P^{Y_i}	$\theta = H$	$\theta = L$
$s^{Y_i} = 1$	$1 - \sigma'_i(h^t) \pi_i^H(h^t)$	$1 - \sigma'_i(h^t) \pi_i^L(h^t)$
$s^{Y_i} = 0$	$\sigma'_i(h^t) \pi_i^H(h^t)$	$\sigma'_i(h^t) \pi_i^L(h^t)$

with $\sigma'_i(h^t) > \sigma_i(h^t)$. Then we can write

$$P^{X_i} = \Phi P^{Y_i},$$

where the stochastic matrix Φ is given by:

Φ	$s^{Y_i} = 1$	$s^{Y_i} = 0$
$s^{X_i} = 1$	1	$\frac{\sigma'_i(h^t) - \sigma_i(h^t)}{\sigma'_i(h^t)}$
$s^{X_i} = 0$	0	$\frac{\sigma_i(h^t)}{\sigma'_i(h^t)}$

Since Φ is a stochastic matrix that is independent of θ , X_i is a garbling of Y_i , and therefore Y_i is sufficient for X_i .

Since the individual exit decisions X_i are independent (conditional on the informational status of the players), the same argument as above applies for the joint experiments $X := \times_{i=1}^{n(h^t)} X_i$ and $Y = \times_{i=1}^{n(h^t)} Y_i$.

Finally, consider two experiments $X^\omega = (X, Z)$ and $Y^\omega = (Y, Z)$ on $\Omega = \{G, B\}$ where X and Y are as above and Z is an experiment with outcomes in $S^Z = \{0, 1\}$. Since θ is correlated with ω , the information contained in X and Y is also information on Ω . We interpret Z as the individual learning experiment on own type and hence the matrix of conditional probabilities for that experiment is given by P^Z :

P^Z	$\omega = G$	$\omega = B$
$s^Z = 1$	$\lambda\Delta$	0
$s^Z = 0$	$1 - \lambda\Delta$	1

Since (X, Z) is a garbling of (Y, Z) by the argument above, we know that (Y, Z) is sufficient for (X, Z) with respect to Ω . The assertion that $P_i^{t+1}(h^t, \sigma(h^t))$ second order stochastically dominates $P_i^{t+1}(h^t, \sigma'(h^t))$ follows from Blackwell's theorem. ■

Proof of Lemma 1. i) Assume that σ is a symmetric equilibrium profile and $p(h^t) > p^*(\Delta)$. If i continues for one period and then exits, her payoff is:

$$p(h^t) \lambda \Delta (v + \delta V^+(\Delta)) - c \Delta > 0,$$

where the inequality follows from (15). This implies that $\sigma(h^t) = 0$ if $p(h^t) > p^*(\Delta)$.

ii) Assume that $p(h^t) < p^*(\Delta)$. Note from (15) that for such a belief we have:

$$p(h^t) \lambda (v + \delta V^+(\Delta)) < c. \tag{27}$$

We show that $\sigma(h^t) > 0$. Assume on the contrary that $\sigma(h^t) = 0$. Let h^τ be the first continuation history of h^t on the equilibrium path such that $\sigma(h^\tau) > 0$. In other words, $h^\tau = (h^t, a^t, \dots, a^{\tau-1})$ where $a^s = (1, \dots, 1)$ and $\sigma(h^s) = 0$ for all $s \in \{t, t+1, \dots, \tau-1\}$, and $\sigma(h^\tau) > 0$.

First, note that if no such h^τ exists, then the players never exit, and their value calculated at history h^t is given by:

$$V(h^t) = \frac{p(h^t) \lambda v \Delta - c \Delta}{1 - \delta} < 0,$$

where the inequality follows from (27). This is a contradiction, because equilibrium value must be at least zero due to the exit option.

Since exiting is in the support of the equilibrium strategy at h^τ , we have

$$V(h^\tau) = 0,$$

and hence the value in period $\tau - 1$ is given by

$$V(h^{\tau-1}) = -c\Delta + p(h^{\tau-1}) \lambda \Delta (v + \delta V^+(\Delta)) < 0,$$

which contradicts optimal behavior at $h^{\tau-1}$. It follows that $\sigma(h^t) > 0$ if $p(h^t) < p^*(\Delta)$. This implies that $V(h^t) = 0$ if $p(h^t) < p^*(\Delta)$. If $p(h^t) = p^*(\Delta)$, we have

$$V(h^t) = \max \{0, p^*(\Delta) \lambda \Delta (v + \delta V^+(\Delta)) - c\Delta\} = 0.$$

iii) Since $V(h^t) = 0$ whenever $p(h^t) \leq p^*$, and since $\sigma(h^t) = 0$ for $p(h^t) > p^*$, the pure strategy $a_m(p(h^t))$ is a best response for each player after each history h^t , given the strategy profile σ . For any h^t such that $p(h^t) > p^*(\Delta)$, $p(h^{t+1})$ is updated on the equilibrium path using the isolated player's belief updating formula in equation (2) since $\sigma(h^t) = 0$. Therefore, $V(h^t) = V_m(p(h^t))$ for each h^t . ■

Proof of Lemma 2. By Lemma 1,

$$\mathbb{E}V(h^{t+1}) = \mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t))).$$

Furthermore, $V_m(p)$ is convex in p and $P^{t+1}(h^t; \sigma(h^t))$ is second order stochastically decreasing in $\sigma(h^t)$ by Proposition 2 and hence the first claim follows.

To prove the second claim, suppose that $p(h^t) < p^*(\Delta)$ and that there exists a $\sigma(h^t)$ such that

$$\delta \mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t))) = c(h^t) \Delta. \quad (28)$$

We claim that for all $\sigma'(h^t) > \sigma(h^t)$,

$$\mathbb{E}V_m(P^{t+1}(h^t; \sigma'(h^t))) > c(h^t) \Delta. \quad (29)$$

To see this, consider the exit decision of player i when all players use the symmetric strategy $\sigma(h^t)$ and (28) holds. Since $p(h^t) < p^*(\Delta)$, there must be an exit decision vector \widehat{a}_{-i}^t for players other than i that makes player i 's exit decision a_i^t pivotal in the following sense. In what follows, we use the notation $p(h^t; \sigma(h^t))$ and $V(h^t; \sigma(h^t))$ to emphasize the dependence of these quantities on the equilibrium strategies. For $h_0^{t+1} := (h^t, (\widehat{a}_{-i}^t, 0))$,

$$p(h_0^{t+1}; \sigma(h^t)) < p^*(\Delta), \quad (30)$$

and for $h_1^{t+1} := (h^t, (\widehat{a}_{-i}^t, 1))$,

$$p(h_1^{t+1}; \sigma(h^t)) > p^*(\Delta). \quad (31)$$

Furthermore,

$$\Pr\{A_{-i}(h^t) = \widehat{a}_{-i}^t\} > 0.$$

Suppose next that player i exits with probability $\sigma'_i > \sigma(h^t)$ and all other players exit with probability $\sigma(h^t)$ after history h^t . We consider the beliefs of an arbitrary player $j \neq i$ following this change in the strategy profile at history h^t .

Denote the profile where all players but i exit with probability $\sigma(h^t)$ and i exits with probability σ'_i by (σ_{-i}, σ'_i) . By Proposition 2, and by the convexity of $V_m(p)$, we know that for every a_{-i}^t

$$\delta \mathbb{E}_{a_i} V((h^t, (a_{-i}^t, a_i)); (\sigma_{-i}, \sigma'_i)) \geq \delta \mathbb{E}_{a_i} V((h^t, (a_{-i}^t, a_i)); \sigma(h^t)).$$

Therefore, the payoff of players other than i is strictly increasing in σ'_i , if for the exit vector \widehat{a}_{-i}^t , the previous inequality is strict, i.e.:

$$\delta \mathbb{E}_{a_i} V((h^t, (\widehat{a}_{-i}^t, a_i)); (\sigma_{-i}, \sigma'_i)) > \delta \mathbb{E}_{a_i} V((h^t, (\widehat{a}_{-i}^t, a_i)); \sigma(h^t)).$$

But this follows immediately from (30) - (31) and the strict convexity of the isolated player value function V_m in the neighbourhood of $p^*(\Delta)$. To see this, note that

$$p(h_0^{t+1}; (\sigma_{-i}, \sigma'_i)) = p(h_0^{t+1}; \sigma(h^t)) < p^*(\Delta),$$

$$p(h_1^{t+1}; (\sigma_{-i}, \sigma'_i)) > p(h_1^{t+1}; \sigma(h^t)) > p^*(\Delta),$$

and:

$$0 = \frac{\partial^- V_m(p(h_0^{t+1}; \sigma(h^t)))}{\partial p} < \frac{\partial^- V_m(p(h_1^{t+1}; \sigma(h^t)))}{\partial p},$$

where $\frac{\partial^- V_m(p)}{\partial p}$ denotes the derivative from the left (which exists by the convexity of $V_m(p)$) of V_m at p .

Starting with the strategy profile (σ_{-i}, σ'_i) , change the exit probability of all players $j \neq i$ to $\sigma'_j = \sigma'_i$, and denote the resulting symmetric profile by $\sigma'(h^t)$. By Proposition 2, the payoff to all players is weakly increased. Therefore for all j ,

$$\begin{aligned} \mathbb{E}V_j(h^{t+1}; \sigma'(h^t)) &= \mathbb{E}V_m(P^{t+1}(h^t; (\sigma'_{-i}, \sigma'_i))) \geq \mathbb{E}V_m(P^{t+1}(h^t; (\sigma_{-i}, \sigma'_i))) \\ &> \mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t))) = \mathbb{E}V_j(h^{t+1}; \sigma(h^t)). \end{aligned}$$

■

Proof of Theorem 1. All we have to do is to check that $\sigma(h^t)$ is optimal for all players under all three cases given in the Theorem, and that this is the only symmetric exit probability with this property.

Lemma 1 implies that it is optimal to stay (exit) at h^t iff $\sigma(h^t)$ satisfies

$$\delta\mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t))) \geq (\leq) c(h^t) \Delta. \quad (32)$$

Consider now cases i) - iii) below. These cases cover all possibilities and are mutually exclusive, because $\mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t)))$ is increasing in $\sigma(h^t)$ by Lemma 2.

i) Assume that

$$\delta\mathbb{E}V_m(P^{t+1}(h^t; 0)) > c(h^t) \Delta.$$

Then it is strictly optimal for all the players to stay in the game if $\sigma(h^t) = 0$. Moreover, by Lemma 2, $\delta\mathbb{E}V_m(P^{t+1}(h^t; x)) > c(h^t) \Delta$ for all $x \geq 0$, so $\sigma(h^t) = 0$ is the unique symmetric equilibrium action in that case.

ii) Assume that

$$\delta\mathbb{E}V_m(P^{t+1}(h^t; 0)) \leq c(h^t) \Delta \leq \delta\mathbb{E}V_m(P^{t+1}(h^t; 1)).$$

First note that $\mathbb{E}V_m(P^{t+1}(h^t; \sigma(h^t)))$ is continuous in $\sigma(h^t)$ as a result of the continuity of the Bayes' rule in $\sigma(h^t)$. Lemma 2 implies that there is a unique value $\sigma^*(h^t)$ for which

$$\delta\mathbb{E}V_m(P^{t+1}(h^t; \sigma^*(h^t))) = c(h^t) \Delta.$$

Moreover, for all $\sigma(h^t) < \sigma^*(h^t)$ the strictly optimal action is to exit, and for all $\sigma(h^t) > \sigma^*(h^t)$ the strictly optimal action is to stay. Thus, $\sigma^*(h^t)$ is the unique symmetric equilibrium action in this case.

iii) Assume that

$$\delta\mathbb{E}V_m(P^{t+1}(h^t; 1)) < c(h^t) \Delta.$$

Then it is strictly optimal for all the players to exit if $\sigma(h^t) = 1$. Moreover, by Lemma 2, $\delta\mathbb{E}V_m(P^{t+1}(h^t; x)) < c(h^t) \Delta$ for all $x \leq 1$, so $\sigma(h^t) = 1$ is the unique symmetric equilibrium action in that case. ■

Proof of Lemma 3. i) We claim first that for each $\varepsilon > 0$ there exists a $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$,

$$p(h^t) < p^*(\Delta) + \varepsilon \text{ for all } h^t \text{ such that } t \geq t^*(\Delta).$$

Fix an $\varepsilon > 0$ and suppose that the claim does not hold. Then, with an arbitrarily small Δ , we must be able to find some history h^t such that if there are no exits, the belief jumps above $p^*(\Delta) + \varepsilon$:

$$p(h^t, \mathbf{1}) > p^*(\Delta) + \varepsilon, \quad (33)$$

where we denote by $(h^t, \mathbf{1})$ the history at $t + 1$ with no exits at h^t .

Consider the continuation value of an arbitrary uninformed player i at $V(h^t)$. Clearly, this value is bounded from below by the payoff of the event that no player exits at h^t weighted by the probability of that event, which gives us:

$$V(h^t) \geq -c\Delta + \Pr\{A_{-i}^t = \mathbf{1}\}e^{-r\Delta}V_m(p_{t+1}(h^t, \mathbf{1})), \quad (34)$$

where $\Pr\{A_{-i}^t = \mathbf{1}\}$ is the probability that i assigns to the event that no players $j \neq i$ exits at h^t . But since the players are willing to randomize at h^t , we must have $V(h^t) = 0$. Using this fact and (33), we can write (34) as

$$c\Delta > \Pr\{A_{-i}^t = \mathbf{1}\}e^{-r\Delta}V_m(p^*(\Delta) + \varepsilon).$$

Since $V_m(p)$ is strictly increasing for $p > p^*(\Delta)$, there is an $\eta > 0$ such that

$$e^{-r\Delta}V_m(p^*(\Delta) + \varepsilon) > \eta.$$

Hence we have:

$$\Pr\{A_{-i}^t = \mathbf{1}\} < \frac{c\Delta}{e^{-r\Delta}V_m(p^*(\Delta) + \varepsilon)} < \frac{c}{\eta}\Delta. \quad (35)$$

On the other hand, a natural lower bound for $\Pr\{A_{-i}^t = \mathbf{1}\}$ is given by:

$$\Pr\{A_{-i}^t = \mathbf{1}\} \geq \Pr\{\text{all players } j \in \mathcal{N}(h^t) \setminus i \text{ are informed}\}.$$

To evaluate this lower bound, note that for an arbitrary $j \neq i$,

$$\Pr\{j \text{ is informed}\} \geq 1 - \pi_t^L,$$

and therefore,

$$\Pr\{A_{-i}^t = \mathbf{1}\} \geq (1 - \pi_t^L)^{n(h^t)-1}. \quad (36)$$

Since $\pi_t^L < 1$ and $n(h^t) \leq N$, (36) contradicts (35) for small enough Δ , and the claim is established.

It remains to show that if $h^t \in H^f(\Delta)$, then $p(h^t) > p^*(\Delta) - \varepsilon$ for sufficiently small Δ . Note that $h^t \in H^f(\Delta)$ means that $h^t = (h^{t-1}, \mathbf{1})$. If $p(h^{t-1}) < p^*(\Delta)$, then $p(h^{t-1}, \mathbf{1}) > p^*(\Delta)$, otherwise the players would strictly prefer to exit at h^{t-1} . If $p(h^{t-1}) \geq p^*(\Delta)$, then the players stay with probability one and $p(h^{t-1}, \mathbf{1}) > p^*(\Delta)(1 - \lambda\Delta)$ by (2). Hence there exists an $\eta > 0$ such that

$$p(h^t) > p^*(\Delta) - \eta\Delta. \quad (37)$$

Together with the above claim, this establishes part i) of the Lemma.

ii) Let

$$h^t(k) = (h^{t-1}, a^t(k)),$$

where $a^t(k)$ is a vector of exit decisions where exactly k active players exit at history h^{t-1} . By Bayes' rule, we know that

$$\frac{1 - q(h^t(k))}{q(h^t(k))} = \frac{1 - q(h^t(0))}{q(h^t(0))} \left(\frac{\pi^L(h^{t-1})}{\pi^H(h^{t-1})} \right)^k.$$

By Proposition 1, there is an $\eta > 0$ such that $\frac{\pi^L(h^{t-1})}{\pi^H(h^{t-1})} > 1 + \eta$. Therefore, for all $k \geq 1$, there is an η' such that

$$q(h^t(k)) < q(h^t(0)) - \eta'.$$

By equation (12), there exists a $v > 0$ such that

$$p(h^t(k)) < p(h^t(0)) - v.$$

By part i), for all $\varepsilon > 0$, there exists a $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$,

$$p(h^t(0)) < p^* + \varepsilon.$$

Since $\varepsilon > 0$ can be chosen arbitrarily small, the claim follows. ■

Proof of Proposition 5. i) Take an arbitrary $h^t \in H^f(\Delta)$. If $p(h^t) > p^*(\Delta)$, then no player wants to exit and we have $P_L(X(h^t) > 0) = 0$. Suppose therefore that $p(h^t) \leq p^*(\Delta)$. Since $h^t \in H^f(\Delta)$, we know from (37) in the proof of the previous Lemma that

$$p(h^t) > p^*(\Delta) - \eta\Delta \quad (38)$$

for some η .

Since $p(h^t) \leq p^*(\Delta)$ and $h^t \in H^f(\Delta)$, the players must be indifferent between staying and exiting. In the absence of observational learning, the loss from staying in the game

for an additional period to an uninformed player is

$$\begin{aligned}
& c\Delta - \lambda\Delta p(h^t) (v + e^{-r\Delta}V^+(\Delta)) \\
&= \Delta [c - \lambda p(h^t) (v + e^{-r\Delta}V^+(\Delta))] \\
&< \Delta [c - \lambda (p^*(\Delta) - \eta\Delta) (v + e^{-r\Delta}V^+(\Delta))] \\
&= \Delta\lambda\eta\Delta (v + e^{-r\Delta}V^+(\Delta)) := \widehat{\kappa}(\Delta)^2,
\end{aligned} \tag{39}$$

where the inequality uses (38) and the last line uses (15). Hence the loss is at most quadratic in Δ . This loss must be compensated by the gain from observational learning, which by the same logic as in part i) of the previous Lemma, is bounded from below by

$$\Pr\{A_{-i}^t = \mathbf{1}\}e^{-r\Delta}V_m(p(h_n^t, \mathbf{1})), \tag{40}$$

where $p(h^t, \mathbf{1})$ is the belief at $t+1$ with no exits at h^t , and $\Pr\{A^t = \mathbf{1}\}$ is the probability of no exits at h^t as estimated by an arbitrary uninformed player. A lower bound for the value at $p(h^t, \mathbf{1})$ is given by:

$$V_m(p(h^t, \mathbf{1})) \geq (p(h^t, \mathbf{1}) - p^*(\Delta)) \lambda\Delta (v + e^{-r\Delta}V^+(\Delta)). \tag{41}$$

Combining (39), (40), and (41), we have

$$p(h^t, \mathbf{1}) - p^*(\Delta) < \frac{\widehat{\kappa}(\Delta)^2}{\lambda\Delta (v + e^{-r\Delta}V^+(\Delta)) \Pr\{A_{-i}^t = \mathbf{1}\}e^{-r\Delta}}.$$

Noting that $\Pr\{A_{-i}^t = \mathbf{1}\}$ is bounded away from 0 for the same reason as in the proof of Lemma 3, we can conclude that there is a $\eta' > 0$ such that

$$p(h^t, \mathbf{1}) < p^*(\Delta) + \eta'\Delta. \tag{42}$$

Combining (38) and (42), we can conclude that

$$\frac{p(h^t, \mathbf{1})}{1 - p(h^t, \mathbf{1})} \cdot \frac{1 - p(h^t)}{p(h^t)} < \frac{p^*(\Delta) + \eta'\Delta}{1 - p^*(\Delta) - \eta'\Delta} \cdot \frac{1 - p^*(\Delta) + \eta\Delta}{p^*(\Delta) - \eta\Delta} < 1 + \eta''\Delta \tag{43}$$

for some $\eta'' > 0$.

Let $\xi^\theta(h^t)$ denote the probability with which an arbitrary player exits at h^t , conditional on θ . Bayes' rule gives:

$$\frac{p(h^t, \mathbf{1})}{1 - p(h^t, \mathbf{1})} = \frac{(1 - \xi^H(h^t))^{n(h^t)-1} p(h^t)}{(1 - \xi^L(h^t))^{n(h^t)-1} (1 - p(h^t))} > \frac{1 - \xi^H(h^t)}{1 - \xi^L(h^t)} \frac{p(h^t)}{1 - p(h^t)}. \tag{44}$$

By Proposition 1, there is some $\alpha < 1$ such that

$$\xi^H(h^t) < \alpha\xi^L(h^t),$$

and therefore by (44),

$$\frac{p(h^t, \mathbf{1})}{1 - p(h^t, \mathbf{1})} \cdot \frac{1 - p(h^t)}{p(h^t)} > \frac{1 - \alpha \xi^L(h^t)}{1 - \xi^L(h^t)}. \quad (45)$$

Combining (43) and (45) gives:

$$\xi^L(h_n^t) < \frac{\eta''}{2(1 - \alpha)} \Delta.$$

Therefore,

$$P_L(X(h^t) > 0) < N \xi^L(h^t) < \frac{N \eta''}{2(1 - \alpha)} \Delta := \kappa \Delta.$$

ii) From part ii) of Lemma 3, we know that there is a $\bar{\Delta} > 0$ and $v > 0$ such that $p(h^t) < p^*(\Delta) - v$ for all $h^t \in H^w(\Delta)$, $\Delta < \bar{\Delta}$. Indifference requires that $p(h^{t+1}, \mathbf{1}) > p^*(\Delta)$. The result follows then from Bayes' rule. ■

Proof of Proposition 6. By Proposition 5, the probability of exit per period in flow mode is at most proportional Δ . Therefore, we know that in the limit $\Delta \rightarrow 0$, there is a well defined hazard rate for the exits. Moreover, by (1) the probability of exit for an individual player is proportional to the probability with which that player is uninformed, and therefore in the limit $N \rightarrow \infty$, the hazard rate of exit conditional on state must be proportional to $\pi^\theta(\tau)$:

$$\frac{f_{k+1}^\theta(\tau | \tau_1, \dots, \tau_k)}{1 - F_{k+1}^\theta(\tau | \tau_1, \dots, \tau_k)} = \pi^\theta(\tau) K,$$

where K is a number that we determine using Lemma 3.

Consider the change in beliefs within short $d\tau$ induced by a given exit intensity K . Denote by $q(\tau)$ the belief of an uninformed player on the aggregate state at real time instant τ . Within $d\tau$, this belief changes for two reasons. First, a good signal arrives with probability $p^\theta(\tau) \lambda d\tau$, and second, an exit takes place with probability $\pi^\theta(\tau) K d\tau$. Therefore, Bayes' rule gives:

$$q(\tau + d\tau) = \frac{q(\tau) (1 - p^H(\tau) \lambda d\tau) (1 - \pi^H(\tau) K d\tau)}{q(\tau) (1 - p^H(\tau) \lambda d\tau) (1 - \pi^H(\tau) K d\tau) + (1 - q(\tau)) q(\tau) (1 - p^L(\tau) \lambda d\tau) (1 - \pi^L(\tau) K d\tau)}.$$

With $d\tau$ small, this can be expressed as:

$$\frac{dq(\tau)}{d\tau} = q(\tau) (1 - q(\tau)) [(\pi^L(\tau) - \pi^H(\tau)) K + (p^L(\tau) - p^H(\tau)) \lambda]. \quad (46)$$

On the other hand, we know by Lemma 3 that when $\Delta \rightarrow 0$, the belief of an uninformed player on aggregate state must be given by

$$q(\tau) = \frac{p^* - p^L(\tau)}{p^H(\tau) - p^L(\tau)}. \quad (47)$$

Differentiating this with respect to τ gives us:

$$\frac{dq(\tau)}{d\tau} = \lambda \frac{p^H(\tau)(1-p^H(\lambda))(p^* - p^L(\tau)) + p^L(\tau)(1-p^L(\lambda))(p^H(\tau) - p^*)}{(p^H(\tau) - p^L(\tau))^2}, \quad (48)$$

where we have used:

$$\frac{dp^\theta(\tau)}{d\tau} = -\lambda p^\theta(\tau)(1-p^\theta(\tau)).$$

Equating (46) and (48), using (47), and solving for K gives us:

$$K = \lambda \frac{p^*(1-p^*)(p^H(\tau) - p^L(\tau))}{(p^* - p^L(\tau))(p^H(\tau) - p^*)(\pi^L(\tau) - \pi^H(\tau))}.$$

■

Proof of Proposition 7. Consider an arbitrary period of an exit wave, where belief is given by q_s . The number of exits at this period is binomially distributed with parameters $\sigma_n \pi^\theta(\theta)$ and n , where n denotes the number of active players in the game and σ_n denotes the exit probability of an individual uninformed player. In case there is no exit, the next period belief is given by Bayes' rule by:

$$q_{s+1}(0) = \frac{(1 - \sigma_n \pi^H(\tau))^{n-1} q_s}{(1 - \sigma_n \pi^H(\tau))^{n-1} q_s + (1 - \sigma_n \pi^L(\tau))^{n-1} (1 - q_s)}.$$

Lemma 3 requires that

$$\lim_{n \rightarrow \infty} \frac{(1 - \sigma_n \pi^H(\tau))^{n-1} q_s}{(1 - \sigma_n \pi^H(\tau))^{n-1} q_s + (1 - \sigma_n \pi^L(\tau))^{n-1} (1 - q_s)} = q^*(\tau),$$

which gives us:

$$\frac{q^*(\tau)}{1 - q^*(\tau)} = \frac{q_s}{1 - q_s} \lim_{n \rightarrow \infty} \frac{(1 - \sigma_n \pi^H(\tau))^{n-1}}{(1 - \sigma_n \pi^L(\tau))^{n-1}}.$$

Taking logarithm on both sides and evaluating limits gives us:

$$\lim_{n \rightarrow \infty} (n-1) \sigma_n = \frac{1}{\pi^L(\tau) - \pi^H(\tau)} \log \left(\frac{q^*(\tau)}{1 - q^*(\tau)} \frac{1 - q_s}{q_s} \right).$$

Therefore,

$$\lim_{n \rightarrow \infty} n \sigma_n \pi^\theta(\theta) = \frac{\pi^\theta}{\pi^L(\tau) - \pi^H(\tau)} \log \left(\frac{q^*(\tau)}{1 - q^*(\tau)} \frac{1 - q_s}{q_s} \right),$$

and the result is given by the Poisson approximation of Binomial distribution.

The second claim is an immediate consequence of the Bayes' rule with k exits. ■

Appendix B: Proofs for Section 5

This appendix contains the proofs for Proposition 3, Proposition 4, and Theorem 2. Our arguments rely on the convergence of the outside observer's belief $\widehat{q}(h^t)$, which we define as the posterior belief of the event $\{\theta = H\}$ at public history h^t :

$$\widehat{q}(h^t) = \frac{q_0 P_H\{h^t\}}{q_0 P_H\{h^t\} + (1 - q_0) P_L\{h^t\}}.$$

Note that $\widehat{q}(h^t)$ differs slightly from $\widehat{q}_i(h^t)$ that was defined in Section 3.2: $\widehat{q}(h^t)$ is based on actions of all players (and represents therefore a true *outside* observer), while $\widehat{q}_i(h^t)$ is based on actions of players other than i .

To link $\widehat{q}(h^t)$ to the beliefs of actual players in the game, note that player i 's belief $q_i(h^t)$ differs from $\widehat{q}(h^t)$ only to the extent that i 's private information affects her belief. Lemma 4 below guarantees that player i 's private history cannot overwhelm a sufficiently strong public history.¹⁸

Lemma 4 *Suppose that $\rho^H < 1$. Then for all $\varepsilon > 0$, there is some $\delta > 0$ such that the following implications hold for all i :*

$$\widehat{q}(h^t) \geq 1 - \delta \implies q_i(h^t) \geq 1 - \varepsilon \text{ and} \quad (49)$$

$$\widehat{q}(h^t) \leq \delta \implies q_i(h^t) \leq \varepsilon. \quad (50)$$

Proof. Recall from Section 3.2 the definition of $\widehat{q}_i(h^t)$ as the belief based on public histories of all players other than i . In addition to this, $\widehat{q}(h^t)$ also conditions on actions of i . Consider the effect of this additional information. The most favorable piece of evidence in terms of state $\theta = H$ that could ever be obtained from i 's actions is the one that fully reveals i to be informed. The likelihood ratio of being informed across the states is $(1 - \pi_t^H) / (1 - \pi_t^L)$, so Bayesian rule gives us an upper bound for $\widehat{q}(h^t)$ as expressed in terms of $\widehat{q}_i(h^t)$:

$$\frac{\widehat{q}(h^t)}{1 - \widehat{q}(h^t)} \leq \frac{1 - \pi_t^H}{1 - \pi_t^L} \frac{\widehat{q}_i(h^t)}{1 - \widehat{q}_i(h^t)}. \quad (51)$$

On the other hand, we can write the relationship between $q_i(h^t)$ and $\widehat{q}_i(h^t)$ using (11):

$$\frac{q_i(h^t)}{1 - q_i(h^t)} = \frac{\pi_t^H}{\pi_t^L} \frac{\widehat{q}_i(h^t)}{1 - \widehat{q}_i(h^t)}. \quad (52)$$

¹⁸In the pure common values case, where $\rho^H = 1$, the ratio $\frac{\pi_t^H}{\pi_t^L} \rightarrow 0$ as $t \rightarrow \infty$. In that case the statement below holds for all t up to an arbitrary, fixed \bar{t} . This modification is not essential for any of our results.

Combining (51) and (52) gives us:

$$\frac{q_i(h^t)}{1 - q_i(h^t)} \geq \frac{\pi_t^H}{\pi_t^L} \frac{1 - \pi_t^L}{1 - \pi_t^H} \frac{\widehat{q}(h^t)}{1 - \widehat{q}(h^t)}. \quad (53)$$

By (10), π_t^H and $1 - \pi_t^L$ are bounded away from zero, and therefore the first equation of Lemma 4 follows directly from (53). The second equation follows from the fact that $q_i(h^t) \leq \widehat{q}(h^t)$, which in turn follows directly from (10) and (11). ■

Proof of Proposition 3

Our proof strategy is to follow the evolution of the outside observer's belief along a filtration that samples the players' actions sequentially one player at a time. We show that this belief must converge to truth as the number of exits increases. Furthermore this implies the convergence of the actual players' beliefs in the original filtration where all actions within a period are sampled simultaneously. The key step in the argument is Lemma 5 below, which implies that the belief process that we consider has a strong drift towards truth when sampled at the points where any player exits. With this Lemma at hand, the rest of the argument is a relatively straightforward application of Theorem A.1. of Fudenberg & Levine (1992).

We use index $s \in \mathbb{N}$ to track the moments of observation starting from period $t^*(\Delta)$ in the following way. At $s = 1$ the action of player 1 in period $t^*(\Delta)$ is observed. At $s = 2$, the action of player 2 in period $t^*(\Delta)$ is observed, and so on. Once the decisions of all N players in period $t^*(\Delta)$ have been sampled, the process moves to the next time period. At $s = N + 1$ player 1's action in period $t^*(\Delta) + 1$ is observed, and so on. This means that we map every $s \in \mathbb{N}$ to the corresponding period $t(s)$ and player $i(s)$ as follows:

$$\begin{aligned} t(s) &: = \left\lfloor \frac{s}{N} \right\rfloor + t^*(\Delta), \\ i(s) &: = s - N \cdot \left\lfloor \frac{s}{N} \right\rfloor. \end{aligned}$$

Let ξ_s^θ denote the exit probability of player $i(s)$ in period $t(s)$ with equilibrium strategy profile σ (nothing in the proof requires this to be symmetric):

$$\xi_s^\theta := \xi_{i(s)}^\theta(h^{t(s)}),$$

where we set $\xi_i^\theta(h^t) = 0$ if $a_{i(s)}^{t(s)-1} = 0$ (that is, probability of exit is zero for a player that has already exited). We use $x_s \in \{0, 1\}$ as an indicator for player $i(s)$ exiting in period $t(s)$:

$$x_s = \begin{cases} 1 & \text{if } a_{i(s)}^{t(s)-1} = 1 \text{ and } a_{i(s)}^{t(s)} = 0 \\ 0 & \text{otherwise} \end{cases}.$$

We use notation h_s to refer to the history of exits up to s :

$$h_s = \left(a_{i(1)}^{t(1)}, \dots, a_{i(s)}^{t(s)} \right),$$

and we denote by \widehat{q}_s the belief process of the outside observer, who observes the players sequentially:

$$\widehat{q}_s = \Pr \{ \theta = H \mid h_s \}, \quad s \in \mathbb{N}.$$

By Bayes' rule, this belief evolves according to:

$$\begin{aligned} \widehat{q}_0 &= q_0 \\ \widehat{q}_s &= \begin{cases} \frac{\widehat{q}_{s-1} \xi_s^H}{\widehat{q}_{s-1} \xi_s^H + (1 - \widehat{q}_{s-1}) \xi_s^L} & \text{if } x_s = 1 \\ \frac{\widehat{q}_{s-1} (1 - \xi_s^H)}{\widehat{q}_{s-1} (1 - \xi_s^H) + (1 - \widehat{q}_{s-1}) (1 - \xi_s^L)} & \text{if } x_s = 0 \end{cases}, \quad s = 1, 2, \dots \end{aligned} \quad (54)$$

Note that for all $s = t \cdot N$, $t \in \mathbb{N}$, the belief \widehat{q}_s coincides with the belief at period t of an outside observer that observes all players simultaneously at each period:

$$\widehat{q}_s = \widehat{q}(h^t), \quad s = t \cdot N, \quad t \in \mathbb{N}.$$

For all other values of s , \widehat{q}_s is the belief of an outside observer who has observed only a subset of players in the last period.

Let X_∞ denote the total number of players that exit the game:

$$X_\infty := \sum_{t=0}^{\infty} X(h^t).$$

We next define an increasing sequence of natural numbers $\{s(k)\}_{k=1}^{X_\infty}$ as follows:

$$\begin{aligned} s(0) &= 0, \\ s(k) &= \min \{ s > s(k-1) \mid x_s = 1 \}, \quad k = 1, \dots, X_\infty. \end{aligned}$$

Hence, $\{\widehat{q}_{s(k)}\}_{k=1}^{X_\infty}$ is a subset of $\{\widehat{q}_s\}_{s=1}^{\infty}$ sequence, that samples the beliefs immediately after realized exits.

Define:

$$L_k := \begin{cases} \frac{1 - \widehat{q}_{s(k)}}{\widehat{q}_{s(k)}} & \text{for } k = 1, \dots, X_\infty. \\ 0 & \text{for } k = X_\infty + 1, \dots \end{cases} \quad (55)$$

In words, L_k is the likelihood ratio for the event $\{\theta = L\}$ sampled after realized exits. It is clear that under the event $\{\theta = H\}$, this process is a martingale. The next lemma is the key to our argument, and it states that this process is an *active supermartingale*, as defined in Fudenberg & Levine (1992).

Lemma 5 *There exists an $\eta > 0$ such that*

$$P_H(|L_{k+1}/L_k - 1| > \eta | h_{s(k)}) > \eta \quad (56)$$

for all $L_k > 0$.

Proof. Note first that

$$\{|L_{k+1}/L_k - 1| \leq \eta\} \iff (1 - \eta)L_k \leq L_{k+1} \leq (1 + \eta)L_k. \quad (57)$$

By Proposition 1 and (13), there is some $\gamma > 0$ such that

$$\frac{\xi_s^L}{\xi_s^H} > 1 + \gamma \quad (58)$$

for all $s \in \mathbb{N}$. We choose some $\eta > 0$ small enough to ensure that

$$\frac{(1 + \eta)}{(1 - \eta)^2} < 1 + \gamma. \quad (59)$$

Write

$$\tilde{L}_s := \frac{1 - \hat{q}_s}{\hat{q}_s}, \quad s \in \mathbb{N}.$$

Note that $L_k = \tilde{L}_{s(k)}$ for $k = 1, \dots, X_\infty$. Using (54) and (58), we have:

$$\tilde{L}_s = \begin{cases} \frac{\xi_s^L}{\xi_s^H} \tilde{L}_{s-1} > (1 + \gamma) \tilde{L}_{s-1} & \text{if } x_s = 1 \\ \frac{(1 - \xi_s^L)}{(1 - \xi_s^H)} \tilde{L}_{s-1} < \tilde{L}_{s-1} & \text{if } x_s = 0 \end{cases}$$

By definition of $s(k)$, we have $x_{s(k+1)} = 1$, and therefore, we have

$$\tilde{L}_{s(k+1)} > (1 + \gamma) \tilde{L}_{s(k+1)-1}. \quad (60)$$

Noting that $L_{k+1} = \tilde{L}_{s(k+1)}$ and $L_k = \tilde{L}_{s(k)}$, and using (57) and (60), we have:

$$\{|L_{k+1}/L_k - 1| \leq \eta\} \implies \left\{ \tilde{L}_{s(k+1)-1} < \frac{1 + \eta}{1 + \gamma} \tilde{L}_{s(k)} \right\}. \quad (61)$$

Let \bar{s} be the first observation point after $s(k)$ at which $\tilde{L}_{\bar{s}}$ is below $\frac{1 + \eta}{1 + \gamma} \tilde{L}_{s(k)}$ in case there are no exits:

$$\bar{s} := \min \left\{ s' > s(k) : \left(\prod_{j=s(k)+1}^{s'} \frac{1 - \xi_j^L}{1 - \xi_j^H} \right) \tilde{L}_{s(k)} < \frac{1 + \eta}{1 + \gamma} \tilde{L}_{s(k)} \right\}.$$

Then it follows from (61) and (57) that:

$$\begin{aligned} & \{|L_{k+1}/L_k - 1| \leq \eta\} \\ \implies & \left\{ x_s = 0 \forall s = s(k) + 1, \dots, \bar{s} \text{ and } \tilde{L}_{s(k+1)} > \frac{1-\eta}{1+\eta} (1+\gamma) \tilde{L}_{\bar{s}} \right\}. \end{aligned} \quad (62)$$

But, since \tilde{L}_s is a supermartingale under $\theta = H$, we have

$$\mathbb{E} \left(\tilde{L}_{s(k+1)} \mid h^{\bar{s}}, \theta = H \right) < \tilde{L}_{\bar{s}},$$

which implies the following (using the fact that $\tilde{L}_{s(k+1)}$ is bounded from below by 0):

$$P_H \left(\tilde{L}_{s(k+1)} < \frac{1-\eta}{1+\eta} (1+\gamma) \tilde{L}_{\bar{s}} \mid h^{\bar{s}} \right) \geq 1 - \frac{(1+\eta)}{(1-\eta)(1+\gamma)} > \eta,$$

where the last inequality follows from (59). Combining this with (62), we note that

$$P_H (|L_{k+1}/L_k - 1| > \eta \mid h_{s(k)}) > \eta.$$

■

Lemma 5 says that L_k , $k \in \mathbb{N}$ is an *active supermartingale* with activity η , 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 3 (Fudenberg and Levine) *Let $l_0 > 0$, $\varepsilon > 0$, and $\eta \in (0, 1)$ 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 η .

With these preliminaries at hand, we are ready to finish the proof of Proposition 3:

Proof of Proposition 3. Fix an $\varepsilon > 0$. Consider the stochastic process L_k , $k \in \mathbb{N}$, defined in (55). Note from (55) that:

$$L_k \leq \underline{L} \iff \left(\left\{ \hat{q}_{s(k)} \geq \frac{1}{1+\underline{L}} \right\} \text{ or } \{k > X_\infty\} \right).$$

We set \underline{L} small enough to guarantee:

$$L_k \leq \underline{L} \implies \left(\{\hat{q}_{s(k)} > 1 - \varepsilon\} \text{ or } \{k > X_\infty\} \right). \quad (63)$$

By Lemma 5, we know that L_k is an active supermartingale with activity η . By Theorem 3, we can therefore set K high enough to guarantee that

$$P_H \left\{ h^\infty : \sup_{k > K} L_k \leq \underline{L} \right\} \geq 1 - \varepsilon. \quad (64)$$

Combining this with (63), we have

$$P_H\{h^\infty : n(h^t) \leq N - K \text{ and } \widehat{q}(h^t) < 1 - \varepsilon \text{ for some } h^t \in h^\infty\} < \varepsilon. \quad (65)$$

We have now proved the Proposition as regards equation (16). Knowing this, the part concerning equation (17) follows from Bayes' rule as follows. Define the following event:

$$A(K, \varepsilon) := \{h^\infty : n(h^t) < N - K \text{ and } \varepsilon < \widehat{q}(h^t) < 1 - \varepsilon \text{ for some } h^t \in h^\infty\}.$$

Then, by the definition of $A(K, \varepsilon)$, the posterior of $\{\theta = H\}$ conditional on reaching $A(K, \varepsilon)$ must be between ε and $1 - \varepsilon$:

$$\varepsilon < \frac{q_0 P_H(A(K, \varepsilon))}{q_0 P_H(A(K, \varepsilon)) + (1 - q) P_L(A(K, \varepsilon))} < 1 - \varepsilon. \quad (66)$$

Since (65) holds for any ε given large enough K , we know that $P_H(A(K, \varepsilon))$ can be made arbitrarily small by increasing K . Therefore, for (66) to hold, also $P_L(A(K, \varepsilon))$ must go to zero as K is increased, which implies that for any $\varepsilon > 0$, we can find K large enough to ensure that

$$P_L\{h^\infty : n(h^t) \leq N - K \text{ and } \widehat{q}(h^t) < \varepsilon \text{ for some } h^t \in h^\infty\} < \varepsilon.$$

■

Proof of Proposition 4

We work through a number of lemmas. First, we formalize the intuitive fact that whenever the probability that a large number of players exit within the current period is non-negligible, the realized actions generate a precise signal about the state of the world. In particular, if the true state is $\theta = H$, then the beliefs of all players must be very close to one after that period:

Lemma 6 *For all $\varepsilon > 0$ and $q > 0$, there is some $K \in \mathbb{N}$ such that*

$$P_L(X(h^t) > K) > \frac{1}{2} \implies P_H(q(h^{t+1}) > 1 - \varepsilon) > 1 - \varepsilon,$$

whenever $q(h^t) > q$ and $t \geq t^(\Delta)$.*

Proof. Denote

$$\mu_\theta := \mathbb{E}[X(h^t) | \theta] = n(h^t) \xi^\theta(h^t).$$

Since $X(h^t)$ is a random variable that can only take positive values, the following must hold:

$$P_L(X(h^t) > K) > \frac{1}{2} \implies \mu_L > \frac{1}{2}K. \quad (67)$$

By Proposition 1, we know that there is some $\gamma > 0$ such that

$$\frac{\xi^L(h^t)}{\xi^H(h^t)} > 1 + \gamma$$

for all $t \geq t^*(\Delta)$. Consider the random variable

$$Z(h^t) := \frac{X(h^t)}{\mu_L}.$$

We have:

$$\mathbb{E}[Z(h^t) | \theta = H] = \frac{n(h^t) \xi^H(h^t)}{n(h^t) \xi^L(h^t)} < \frac{1}{1 + \gamma}, \quad (68)$$

$$\mathbb{E}[Z(h^t) | \theta = L] = \frac{n(h^t) \xi^L(h^t)}{n(h^t) \xi^L(h^t)} = 1, \quad (69)$$

$$\text{var}[Z(h^t) | \theta = H] = \frac{n(h^t) \xi^H(h^t) (1 - \xi^H(h^t))}{(n(h^t) \xi^L(h^t))^2} < \frac{1}{n(h^t) \xi^L(h^t)} = \frac{1}{\mu_L}, \quad (70)$$

$$\text{var}[Z(h^t) | \theta = L] = \frac{n(h^t) \xi^L(h^t) (1 - \xi^L(h^t))}{(n(h^t) \xi^L(h^t))^2} < \frac{1}{n(h^t) \xi^L(h^t)} = \frac{1}{\mu_L}. \quad (71)$$

Consider the event

$$A = \left(Z(h^t) \leq \frac{1 + \frac{1}{2}\gamma}{1 + \gamma} \right).$$

The formulas (68) - (71) imply that

$$\lim_{\mu_L \rightarrow \infty} P_H(A) = 1 \text{ and } \lim_{\mu_L \rightarrow \infty} P_L(A) = 0.$$

By (67), assuming $P_L(X(h^t) > K) > \frac{1}{2}$ and increasing K will increase μ_L without bound. Hence, the result follows from Bayes' rule by considering the likelihood ratio across states of event A as K is increased. ■

The next step in the proof, Lemma 7, bounds the probability with which a large number of players may exit within an arbitrary single period. By Lemma 6, a random experiment that induces a large number of players to exit with a non-negligible probability is very informative on the aggregate state. Any uninformed player would like to stay in the game until τ_H if she knew the state to be H . Suppose next that the probability of state H is bounded away from zero. As the period length is reduced towards zero, the players would rather wait and observe the result of an informative experiment than exit immediately. Lemma 7 formalizes this argument.

Lemma 7 *For all $\tau < \tau_H$ and $q > 0$, there exist a $K \in \mathbb{N}$ and a $\bar{\Delta} \in \mathbb{R}^+$ such that*

$$q(h^t) > q \implies P_L(X(h^t) > K) < \frac{1}{2}$$

whenever $\Delta < \bar{\Delta}$ and $t \leq t(\tau, \Delta)$.

Proof. Fix a $\tau < \tau_H$ and a $q > 0$. Lemma 6 implies the existence of a function $\phi : \mathbb{N} \rightarrow \mathbb{R}^+$ with

$$\lim_{K \rightarrow \infty} \phi(K) = 0,$$

such that for all h^t , $t \leq t(\tau, \Delta)$, and $q(h^t) > q$:

$$P_L(X(h^t) > K) > \frac{1}{2} \implies P_H(q(h^{t+1}) > 1 - \phi(K)) > 1 - \phi(K). \quad (72)$$

Recall the definition of p_t^H in equation (3), i.e. the belief of a player on her own type conditional on state H . If $\tau < \tau_H$, we can choose an $\eta > 0$ and a $\Delta' > 0$ such that $p_{t(\tau, \Delta)+1}^H > p^*(\Delta) + \eta$ for all $\Delta < \Delta'$. This follows directly from the continuity of $p^*(\Delta)$ and the definition of τ_H . This means that we can choose a K high enough so that

$$q(h^{t+1}) > 1 - \phi(K) \implies p(h^{t+1}) > p^*(\Delta) + \eta. \quad (73)$$

We choose a K such that (72) and (73) hold for all h^t , $t \leq t(\tau, \Delta)$, for which $q(h^t) > q$. Take any such history, and assume that $P_L(X(h^t) > K) > \frac{1}{2}$. Consider next the expected payoff that an uninformed player would get by staying in the game with probability one at that history. We want to find a lower bound for that payoff. Since $q(h^t) > q$, the posterior for $\theta = H$ is bounded from below by q . By (72) and (73), $1 - \phi(K)$ is a lower bound for the probability that $p(h^{t+1}) > p^*(\Delta) + \eta$, conditional on $\theta = H$. Finally, $V_m(p^*(\Delta) + \eta) > 0$ is the value of the isolated player at belief $p^*(\Delta) + \eta$. Therefore, the continuation payoff for a player that stays is bounded from below by:

$$V(h^t) \geq -c\Delta + e^{-r\Delta} \cdot q \cdot (1 - \phi(K)) \cdot V_m(p^*(\Delta) + \eta). \quad (74)$$

We see from (74) that we guarantee $V(h^t) > 0$ by setting Δ small enough and K large enough. Since then it is strictly optimal for any individual player to stay in the game, this contradicts the presumption that $P_L(X(h^t) > K) > \frac{1}{2}$. We thus conclude that for high enough $K \in \mathbb{N}$ and small enough $\bar{\Delta} \in \mathbb{R}^+$ the implication

$$q(h^t) > q \implies P_L(X(h^t) > K) < \frac{1}{2}$$

holds whenever $\Delta < \bar{\Delta}$ and $t \leq t(\tau, \Delta)$. ■

Lemma 8 shows that if a large number of players exit within a period, then the belief of an uninformed player falls to a very low level.

Lemma 8 *For all $\tau < \tau_H$ and $q > 0$, there exist a $K \in \mathbb{N}$ and a $\bar{\Delta} \in \mathbb{R}^+$ such that the following implication holds on the equilibrium path of any game $\Gamma(\Delta, N)$ with $\Delta < \bar{\Delta}$:*

$$\{t \leq t(\tau, \Delta) \wedge q(h^t) > q \wedge X(h^t) > K\} \implies \{q(h^{t+1}) < q\}.$$

Proof. Fix a $\tau < \tau_H$ and a $q > 0$. By Lemma 7, fix $K' \in \mathbb{N}$ and Δ' such that

$$P_L(X(h^t) > K') < \frac{1}{2} \quad (75)$$

whenever $\Delta < \Delta'$, $q(h^t) > q$, $t \leq t(\tau, \Delta)$. Since $\tau < \tau_H$, the same logic that led to (73) allows us to fix $\Delta'' > 0$ and $q^* < 1$ such that whenever $t \leq t(\tau, \Delta)$ and $\Delta < \Delta''$, the following implication holds

$$q(h^t) > q^* \implies p(h^t) > p^*(\Delta). \quad (76)$$

Define $\bar{\Delta} = \min(\Delta', \Delta'')$. For the rest of the proof we assume that $\Delta < \bar{\Delta}$, and we take an arbitrary history h^t such that $t \leq t(\tau, \Delta)$, $q(h^t) > q$, and $\xi^H(h^t) > 0$. Our goal is to find a K such that $X(h^t) > K$ would imply $q(h^{t+1}) < q$.

Consider the expression for the probability of k exits:

$$P_\theta(X(h^t) = k) = \binom{n}{k} (\xi^\theta(h^t))^k (1 - \xi^\theta(h^t))^{n(h^t)-k}. \quad (77)$$

Since $\xi^H(h^t) < \xi^L(h^t)$, it follows by straightforward algebra from (77) that

$$\frac{P_H(X(h^t) = k)}{P_L(X(h^t) = k)} > \frac{P_H(X(h^t) = k')}{P_L(X(h^t) = k')} \text{ for } k < k'. \quad (78)$$

It then also follows that

$$\frac{P_H(X(h^t) = K')}{P_L(X(h^t) = K')} < 2. \quad (79)$$

To see why, assume the contrary. Then, we have

$$P_H(X(h^t) \leq K') = \sum_{k=0}^{K'} P_H(X(h^t) = k) > 2 \cdot \sum_{k=0}^{K'} P_L(X(h^t) = k) > 2 \cdot \frac{1}{2} = 1,$$

where the first inequality uses (78) and the presumption that (79) does not hold, whereas the second inequality follows from (75). But a probability of an event can not be greater than one, so (79) must hold.

Consider next the following expression:

$$\begin{aligned} \frac{P_H(X(h^t) = K' + K'')}{P_L(X(h^t) = K' + K'')} &= \frac{\binom{n}{K' + K''} (\xi^H(h^t))^{K' + K''} (1 - \xi^H(h^t))^{n(h^t) - K' - K''}}{\binom{n}{K' + K''} (\xi^L(h^t))^{K' + K''} (1 - \xi^L(h^t))^{n(h^t) - K' - K''}} \\ &= \left(\frac{\xi^H(h^t)}{\xi^L(h^t)} \right)^{K'} \left(\frac{1 - \xi^H(h^t)}{1 - \xi^L(h^t)} \right)^{n(h^t) - K'} \\ &\quad \cdot \left(\frac{\xi^H(h^t)}{\xi^L(h^t)} \right)^{K''} \left(\frac{1 - \xi^L(h^t)}{1 - \xi^H(h^t)} \right)^{K''} \\ &= \frac{P_H(X(h^t) = K')}{P_L(X(h^t) = K')} \cdot \left(\frac{\xi^H(h^t)}{\xi^L(h^t)} \right)^{K''} \cdot \left(\frac{1 - \xi^L(h^t)}{1 - \xi^H(h^t)} \right)^{K''} \end{aligned}$$

By (79),

$$\frac{P_H(X(h^t) = K')}{P_L(X(h^t) = K')} < 2.$$

Also, since $\xi^L(h^t) > \xi^H(h^t)$, we have

$$\left(\frac{1 - \xi^L(h^t)}{1 - \xi^H(h^t)}\right)^{K''} < 1.$$

By Proposition 1, we have

$$\lim_{K'' \rightarrow \infty} \left(\frac{\xi^H(h^t)}{\xi^L(h^t)}\right)^{K''} = 0,$$

and therefore, we can set K'' high enough to ensure

$$\frac{P_H(X(h^t) = K' + K'')}{P_L(X(h^t) = K' + K'')} < \frac{1 - q^*}{q^*} q. \quad (80)$$

Since $\xi^H(h^t) > 0$, we know from (76) that $q(h^t) < q^*$ (otherwise no player would want to exit). Therefore, Bayes' rule and simple algebra leads to:

$$\begin{aligned} & q(h^{t+1} | X(h^t) = K' + K'') \\ = & \frac{q(h^t) P_H(X(h^t) = K' + K'')}{q(h^t) P_H(X(h^t) = K' + K'') + (1 - q(h^t)) P_L(X(h^t) = K' + K'')} \\ < & \frac{q(h^t) P_H(X(h^t) = K' + K'')}{1 - q(h^t) P_L(X(h^t) = K' + K'')} \\ \leq & \frac{q^* P_H(X(h^t) = K' + K'')}{1 - q^* P_L(X(h^t) = K' + K'')} < q, \end{aligned}$$

where the last inequality follows from (80). By (78), this means that

$$q(h^{t+1} | X(h^t) = k) < q$$

for any $k > K$, where we have set $K := K' + K''$. ■

Finally, we state a lemma that limits the probability with which an outside observer's belief $\hat{q}(h^t)$ could ever get small values if $\theta = H$. This result is simply a formalization of the notion that a Bayesian observer is not likely to be convinced of the untrue state.

Lemma 9 *For all $\varepsilon > 0$, there is a $q > 0$ such that*

$$P_H\{h^\infty : \hat{q}(h^t) \leq q \text{ for some } h^t \in h^\infty\} < \varepsilon.$$

Proof. Consider the event

$$A = \{h^\infty : \hat{q}(h^t) \leq q \text{ for some } h^t \in h^\infty\}.$$

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

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

by the 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 . ■

With Lemmas 8 and 9 at hand, it is now easy to finish the proof of Proposition 4:

Proof of Proposition 4. Fix a $\tau < \tau_H$ and an $\varepsilon > 0$. Using Lemma 9, and noting that the divergence of the uninformed players' belief from outside observer's belief is bounded by Lemma 4, we can choose a $q > 0$ such that

$$P_H \{h^\infty : q(h^t) \leq q \text{ for some } h^t \in h^\infty\} < \varepsilon. \quad (81)$$

Next, by Lemma 8, we can choose a $K \in \mathbb{N}$ and a $\bar{\Delta} \in \mathbb{R}^+$ such that

$$\{t \leq t(\tau, \Delta), q(h^t) > q, X(h^t) > K\} \implies \{q(h^{t+1}) < q\},$$

whenever $\Delta < \bar{\Delta}$. Thus, if there is some $h^t, t \leq t(\tau, \Delta)$, for which $X(h^t) > K$, we must have either $q(h^t) \leq q$ or $q(h^{t+1}) \leq q$. But by (81) this cannot happen with probability greater than ε , and as a result, we have

$$P_H \{h^\infty : X(h^t) > K \text{ for some } t \leq t(\tau, \Delta)\} < \varepsilon$$

if $\Delta < \bar{\Delta}$. ■

Proof of Theorem 2

Proof of Theorem 2. Fix a $\tau < \tau_H$ and an $\varepsilon > 0$. Then, by Lemma 4 and equation (12), we can choose a $q' < 1$ and a $\Delta' > 0$ such that whenever $\Delta < \Delta'$ and $t \leq t(\tau, \Delta)$, the following holds:

$$\hat{q}(h^t) \geq q' \implies p(h^t) > p^*(\Delta). \quad (82)$$

Next, by Proposition 3, we can choose a K' such that

$$P_H \{h^\infty : n(h^t) \leq N - K' \text{ and } \hat{q}(h^t) < q' \text{ for some } h^t \in h^\infty\} < \frac{\varepsilon}{2}. \quad (83)$$

Assume that $\sum_{t=0}^{t(\tau, \Delta)} X(h^t) \geq K'$ (if not, then there is nothing to prove), and denote by $t_{K'}$ the first period with fewer than $N - K'$ active players left in the game:

$$t_{K'} := \min \{t : n(h^t) \leq N - K'\}. \quad (84)$$

Since $\sum_{t=0}^{t(\tau, \Delta)} X(h^t) \geq K'$, we must have

$$t_{K'} \leq t(\tau, \Delta) + 1. \quad (85)$$

Equations (82) and (83) mean that the probability that any player exits in $[t_{K'}, \dots, t(\tau, \Delta)]$ is less than $\frac{\varepsilon}{2}$:

$$P_H \left\{ h^\infty : \sum_{t=t_{K'}}^{t(\tau, \Delta)} X(h^t) > 0 \right\} < \frac{\varepsilon}{2}. \quad (86)$$

By the definition of $t_{K'}$ in (84) we know that

$$\sum_{t=0}^{t_{K'}-2} X(h^t) < K'. \quad (87)$$

Finally, by (85) and Proposition 4, we can find a $\Delta'' \in \mathbb{R}^+$ and a K'' such that

$$P_H \{ h^\infty : X(h^{t_{K'}-1}) > K'' \} < \frac{\varepsilon}{2}. \quad (88)$$

Noting that (86) holds when $\Delta < \Delta'$ and (88) holds when $\Delta < \Delta''$, we may set $K := K' + K''$ and $\bar{\Delta} := \min(\Delta', \Delta'')$, and combine (86) - (88) to get:

$$P_H \left\{ h^\infty : \sum_{t=0}^{t(\tau, \Delta)} X(h^t) > K \right\} < \varepsilon,$$

whenever $\Delta < \bar{\Delta}$. ■

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* 48:477–501.
- Chamley, C. & D. Gale. 1994. "Information Revelation and Strategic Delay in a Model of Investment." *Econometrica* 62:1065–1086.
- Décamps, J.-P. & T. Mariotti. 2004. "Investment Timing and Learning Externalities." *Journal of Economic Theory* 118:80–102.
- 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.
- Rothschild, M. & J. Stiglitz. 1970. "Increasing Risk I: Definition." *Journal of Economic Theory* 2:225–243.
- Smith, L. & P. Sorensen. 2000. "Pathological Outcomes of Observational Learning." *Econometrica* 68:371–398.
- Toxvaerd, F. 2008. "Strategic Merger Waves: A Theory of Musical Chairs." *Journal of Economic Theory* 140:1–26.