Learning from Neighbors in a Changing World

MASTER'S THESIS

by

Ruben van Oosten
(312321)

Department of Economics
Erasmus School of Economics
ERASMUS UNIVERSITY ROTTERDAM

Supervisor: Dr. Bauke Visser

Delft
January 4, 2016

Abstract

This paper analyzes the influence of various network architectures on the ability of a group of people to adapt to a dynamic state of the world. In highly connected societies, interaction with neighbors will clog the information pipeline as individuals overvalue the information generated by the (often incorrect) actions of others, making it impossible to distinguish the signal that contains important information on the true state. As a result, people are unable to change their action over time in a complete network. When people have an unequal number of connections, those with the least number of connections can start the adaptation process and set an example for their neighbors.

Contents

1 Introduction

Information plays a crucial role in any economic decision making process. One of the findings in the social learning literature is that more efficient decision making can be achieved by limiting the information available to those that have to act (Chamley, 2004). Learning from others can be detrimental to society if people decide to mimic others and engage in herd behavior: "everyone doing what everyone else is doing even when their private information suggests doing something quite different." (Banerjee, 1992:798). Collective enthusiasm about a particular decision can severely affect the eventual outcome and multiply gains or losses for those involved.

The goal of this paper is to explain in what way the structure of a network affects the capacity to change for a group of individuals. In doing so, the objective is to isolate cause, the network architecture, from effect, the ability of a connected person to change her action over time in accordance with the dynamic state of the world. Long-term behavior is key, because adaptation is a continuous process that starts after actions have converged. Every individual will have to learn the true state multiple times. The network architecture then causes differences in behavior because it affects the diffusion of information in society. The flow of information is directed by the network characteristics: the number of connections, who is connected to whom and the general level of trust in the correctness of the actions of other people (as a measure of social influence).

To analyze the social learning process I present a model of decision making by a group of myopic individuals. They have to decide simultaneously and repeatedly on one of two possible states of the world. Everyone is aware that the state of the world is dynamic. Hence, the optimal decision changes over time and a different action is required once the state of the world has changed in order to maximize payoffs. Furthermore, no one knows whether any of their previous decisions were correct as payoffs are not visible until the game has ended. To support their decision people have to rely on two sources of information. First, everyone receives a private signal, conditional on the state of the world at that time. The precision of the signal is constant and gives a first indication of what the decision should be. Adapting to the true state is possible for a person without neighbors, because she will always follow her signal. Second, individuals are part of a social network and observe the set of actions taken by their neighbors in the preceding period. The content of this decentralized source of information depends on the location of the agent in the network. The signal precision is common knowledge and the actions of neighbors are judged by a predetermined level of reliability. In conjunction, the precision and reliability determine the weight attached to the two sources of information and their relative influence on the decision. When people observe a lot of neighbors taking the same action and/or neighbors are highly valued, the actions of these neighbors will be influential and can overrule a private signal, also if they are incorrect.

The results show that connected individuals are subject to negative herd externalities. It is shown that a network with more connections decreases the ability of agents to change their action. On the one hand, an increase in the number of connections increases the probability of short-term correct decisions, as more signals (conditional on the underlying state) are inserted in the system. Most agents interpret their private signals correctly and interaction ensures that those that are incorrect will revise their decision. On the other hand, a network with more links between people means that no one is able to change their action over time, because they no longer value the signal as primary source of information. They mostly rely on the actions of their neighbors. A complete network, in which everyone is connected to everyone, gives rise to a persistent herd, where actions converge and none of the people involved is able to deviate. As a consequence, adaptation to a different underlying state is impossible. Consider players A and B deciding whether to invest or not, while observing each others actions. If they both see their neighbor investing time after time, thereby confirming their own belief, after some time they have probably convinced each other and agree that investing is the best option, despite their signal indicating that they are wrong. The cause of the gridlock is that person A is not aware that B has primarily based her decision on the observation of the action of A, and vice versa. If they would have known, and both decided to follow their private signal, they would have behaved differently and retained the ability to adapt. This finding is a consequence of the assumption that there is a predetermined level of trust in the correctness of neighbors' actions, which could be interpreted as peer pressure or a willingness to show conformity with others. Related to the empirical findings on correlation neglect and as a result of bounded rationality, the actions of neighbors are always valued as independent and meaningful. Individuals underestimate the level of correlation of the actions in their neighborhood and an individual with more connections will receive more (often incorrect) information. Treating all information as independent will cause conformity in behavior and prevent adaptation, because everyone underestimates the value of the private signal, which is truly independent and a (noisy) correct reflection of the true state of the world.

Limiting the connections between people can alleviate the tendency to engage in undesirable herd behavior. I show that the same number of agents in an incomplete network might be able to adapt to changing circumstances. These networks exhibit a network dynamic, because time and location determine who can adapt and at what time. Agents in a line network are dependent on the two outer agents to stop the herd. Individuals with one connection receive less information from their network, which will increase the influence of the signal on their decision and direct the belief of the agent towards the true state of the world. When one of these agents decides to change actions others believe that this agent has received valuable new information, which will encourage them to adapt as well. This pattern of behavior generalizes to more complex networks consisting of nodes with an unequal number of connections. Adaptation to a different state will always start at the outskirts and expand inward to highly connected individuals. As long as the individual with the least number of connections can change, she is able to stop the herd and start the adaptation process.

Theories of social learning with a static state of the world have proven to be effective at explaining why actions will converge after some time and why there is long-term enthusiasm about a particular decision (see for example Bikhchandani, Hirshleifer and Welch, 1992; Gale and Kariv, 2003). One of the shortcomings is that these models are less effective at explaining why (groups of) people sometimes decide to radically change course, who should be the first to change and the role of local interaction in this learning process. Roughly speaking, these models are well equipped to explain the inflation of bubbles, but not why they burst. The contribution of this paper is that it describes differences in behavior caused by the way people are connected. A distinguishing feature of the model is that it is able to explain the linkage between two empirical findings: the possibility of an agent to deviate from a herd and the possibility of people to ignore information (a cascade). Goeree, Palfrey, Rogers and McKelvey (2007) search an explanation for the empirical finding that herds are seldom persistent and deviations occur regularly. In a setting of sequential decision making with a fixed state of the world they find that the true state will be revealed over time. If everyone takes the incorrect action, one of the agents is likely to follow a correct signal and deviate. This is the start of an adaptation process to a correct herd for all agents (a self-correcting herd). The drawback of their theory is that it does not allow for agents that ignore information if they achieve a level of certainty over the state of the world. If people would ignore information, no one would be able to follow their signal and change their action. In short, cascades need to be absent for herds to stop. This result is in stark contrast with the empirically proven occurrence of cascades as a type of behavior[1] (Çelen and Kariv, 2004). In

[1] For a model with similar characteristics Çelen and Kariv (2004:497) conclude that: "although cascades are not a theoretical possibility, they are a reality."

the present paper the network architecture determines whether cascades occur and whether herds are persistent. By varying the size and shape of the network we can either observe herds and cascades together, or make cascades absent and herds persistent, or observe very short herds without cascades. In addition, within a network, we can differentiate the level of belief for which people ignore information by varying the number of connections. Moreover, contrary to conventional models of social learning, the convergence of beliefs and actions to a unique state of the world, asymptotic learning, is never optimal. Beliefs do not converge because the optimal action varies with the underlying state and information depreciates. The objective of this thesis is to describe the adaptation process that follows after everyone has decided that mimicking others is optimal and to show which network architectures enable the fashion to become a fad.

As an example, consider a collection of small investors in search of profits and a central bank with its policy to ensure financial stability. Investors aggregate information from several sources. They collect local information from peers by observing what others do and by looking at the volume of traded stock. In addition, the central bank reports on the state of the economy. Every investor will have its own interpretation of these reports, but the information will be considered valuable if everything goes well and there is a balance between the number of investors buying and selling assets. During times of financial turmoil, this trust is fragile. Fear can give rise to sudden changes in collective behavior. As the majority of investors buys or sells assets at the same time, others will be inclined to follow. If there are too many investors taking the same action, all will reassess the value of their information and start to rely on peers as their main source of information. As a consequence, public announcements by the central bank are ineffective and overruled by information generated by the actions of the majority of investors. It is hard to convince people that everything is fine and their investments are safe when they see massive volume of stock being sold by fellow traders. Following the actions taken by peers will then at least ensure constant relative utility. Individually, people make rational decisions, but acting as a group creates a dynamic of its own. A state of groupthink can deepen the crisis and prevents individuals from adequately reacting to information that contradicts their actions. One of the investors must be brave enough to start investing before others will follow and a period of recovery can start. The influence of local interaction on this process, and the conditions that enable the herd to become the primary source of information is the central theme of this paper.

Section 2 presents related literature. Section 3 explains the model. Section 4 describes the restrictions on attainable beliefs as a result of the dynamic nature of the state of the world. Section 5 analyzes differences in behavior for three network structures. Section 6 concludes. Section 7 questions three crucial modelling assumptions. The appendix contains all the proofs.

2 Related Literature

The present paper builds on previous work on social learning of a (dynamic) state of the world and analyzes the influence of a network on the diffusion of information[2].

2.1 Sequential decision making

Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) describe a setting of sequential decision making that can lead to a loss of welfare for society if people are inclined to mimic others, even if their peers have very little information themselves. Individual rational behavior can cause deliberations that seem irrational for an independent bystander, because a lot of valuable information is lost. During a herd, agents underestimate the value of their private signal and overestimate the value of the actions of others.

[2] A broad overview of social learning literature (with empirical applications) is given by Chamley (2004). In addition, Goyal (2007) discusses various aspects of social learning by connected agents.

They think the choices by predecessors are rich in information because the actions are based on information aggregated at that time, but do not realize most of them have been ignoring private information all along. Therefore, the first couple of decisions are crucial to the equilibrium outcome. Smith and Sørensen (2000) emphasize that these results hold because beliefs are bounded as a result of bounded likelihood ratios (bounded away from zero and finite). The informativeness of a binary signal is limited, resulting in a maximum increase or decrease in belief each period. For a prior belief that is sufficiently outspoken (high or low) all available information is discarded in advance, because the signal will not influence the choice of action that period (the start of a cascade). In a model with unbounded beliefs (i.e. unbounded likelihood ratios) signals are arbitrarily strong. In every period there is a potential signal that can change the action. As a result, information is aggregated in the belief parameter and the learning process continues until actions and beliefs converge to the true state (asymptotic learning) (see also Acemoglu, Dahleh, Lobel and Ozdaglar (2011))[3].

The sequential learning literature is related because the agents base their decision on similar sources of information. The interplay between the private signal and the observation of a single predecessor determine how and when herds and cascades arise. Three remedies to counter the tendency to mimic predecessors have been suggested. First, Banerjee (1992) promotes limiting the importance of observations. This forces the agent to make a personal assessment using the private signal and increases ex post welfare. Including a network structure is essentially an extension to this line of thought. The individual has more than one connection and the importance of observations depends on how many neighbors are visible. Second, Bikhchandani et al. (1992) suggest that expanding the information set with public (or external) information can be used as an instrument to delay or correct cascades. In addition, Goeree et al. (2007) add noise

[3]The difference between bounded and unbounded beliefs and the consequences for the learning process are explained in Section 4.2.

to payoffs, thereby assuming that every individual has a different optimal action[4]. The noise is common knowledge, which ensures that each decision maker knows her preference is different from that of others. As a result, there is no common optimal action and everyone has a need for a private signal (see also Goeree, Palfrey and Rogers, 2006). As a consequence, both actions are always chosen with positive probability, ignoring information is no consideration and the resulting information aggregation ensures asymptotic learning. Furthermore, herds are temporary because individuals value their own signal and are likely to act accordingly[5]. Third, as an illustration of why herds are fragile, Bikhchandani et al. (1992) consider the effects of a non-stationary state of the world. A change in the unobservable state variable affects the random signal and reduces the value of information so that a small probability of change can cause differences in behavior. Moscarini et al. (1998) further this thought by showing the depreciation of information resulting from a stochastically changing world. A time-dependent state reduces the information contained in cascades and makes them less valuable so that they end in limited time. Furthermore, a cascade can last longer if the quality of private information decreases or the state change is more predictable, so that the rate of depreciation is lower. Peck and Yang (2011) extend this approach to an n-player setting and apply this framework to explain fluctuation in the business cycle.

A precondition for adaptation to a different state is that beliefs are active. Agents will have to take on information from their signal if they want to adapt to the opposite state or if beliefs need to converge. Goeree et al. (2007) solve this problem by making cascades non-existent. Similarly, Smith and Sørensen (2000) show that unbounded beliefs prevent the tendency to ignore information. In these models all beliefs are always active. In the present paper I follow the approach taken by Moscarini et al.

[4]Their logit quantal response equilibrium model is an extension to the Bayesian standard model by Bikhchandani et al. (1992).

[5]Note that the definition of herds and cascades these authors employ is the exact opposite of the definitions in the present paper (see Section 4.2). They argue that the distinction is not important because cascades never occur (Goeree et al., 2007:739).

(1998) where agents that ignore information will have active beliefs after some time, due to the depreciation of information resulting from a changing state of the world. I replicate their results for a group of connected individuals in Section 4. The main difference is that, even though information has an expiration date, interaction with neighbors might prevent adequate adjustment of actions to the true state when a cascade has ended.

2.2 Simultaneous decision making

Orléan (1995) drops the assumption of sequential decision making and introduces a group opinion, a reference for each agent which includes the information from all observed actions. The main findings are similar to the present model. First, when departing from the sequential context cascades do not have to appear. Second, the weight attached to the observations of others determines the influence of private information on the decision.

Bala and Goyal (1998) study a connected network in which signals are absent but payoffs are observed. As long as one of the agents has an incentive to try a new action there is a positive probability that one of them will find a higher payoff. The result is that beliefs converge and everyone chooses the same action (asymptotic learning). In essence, this is a model of social experimentation, because payoffs of neighbors are directly visible. The difference with a model of social learning is that there is an information externality (others learn from the action) but there is no information asymmetry (payoffs are directly visible) (Gale and Kariv, 2003).

Gale and Kariv (2003) show what happens when (non-strategic) fully rational players can make inferences about the actions of agents that are not observed. Similar to the present paper they find that the results depend on the distribution of connections between agents. Nevertheless, after some time, all agent will choose the same action. In a complete network learning ceases within a few periods, which means there is a high probability of an incorrect action in the long-run. The learning process is spread out over more periods in an incomplete network, which decreases the probability of incorrect actions. The model of Kanoria and Tamuz (2013) is an extension that describes differences in the speed of convergence and calculates the complexity of decisions made by rational Bayesian agents.

Acemoglu et al. (2011) emphasize the role of stochastic generated neighborhood topologies with expanding observations. There is an infinite number of agents observing actions from other neighbors than just a finite subset of agents. If there exist a finite subset, these agents are highly influential and the aggregation of information is troublesome. They show that unbounded private beliefs combined with expanding observations are a necessary and sufficient conditions for asymptotic learning. The result is a consequence of the strong improvement principle, which states that each agent will always be able to imitate the action of her neighbor and will therefore always be at least as well off (comparable to the imitation principle in Bala and Goyal (1998)).

The curse of information principle is a result that has been well established in the literature. Lamberson (2010) finds that the network architecture matters for the diffusion of information. Most notably, depending on prior beliefs and the initial structure of the network, adding links can decrease the efficiency of decision making. Ellison and Fudenberg (1995) find similar results as they show that a society can avoid inefficient herding when (word-of-mouth) communication is limited. González-Avella, Eguíluz, Marsili, Vega-Redondo and San Miguel (2011) describe several network structures and their effect on social learning. An agent follows the signal if a large enough fraction of neighbors are taking this same action. They find that restricted interaction increases the likelihood of asymptotic learning. Lower (average) degree networks (for example with random link probability) lead to more efficient social learning and a higher probability of correct decisions than interaction in a complete network.

One of the downsides of Bayesian learning is that, if the game continues for more than a few periods, agents are required to do complex calculations. Therefore, following DeGroot (1974), a range of models uses simple updating rules. Klumpp (2006) shows that there is a negligible difference in outcomes between Bayesian updating and (simple) linear updating rules[6]. Ellison and Fudenberg (1993) study repeated choices for a group of agents located on a line. They compare rule of thumb decision making, where agents choose the action with the highest payoff in the preceding period, with popularity weighting, where agents take into account the most popular adopted technology in the preceding period. For homogenous players popularity weighting can hurt the short-term convergence of actions in favor of long-term efficiency.

The asymptotic learning outcome and short-term/long-term trade-off is found by most authors. The main difference with these studies is that asymptotic learning is never optimal if the underlying state of the world changes. The objective is not to learn a fixed true state, but to adapt to a different state every so often. Several authors have studied non-Bayesian learning, where agent have to estimate the value of a dynamic state of the world. Frongillo, Schoenebeck and Tamuz (2011) assume a time-varying state following a random walk. The beliefs in complete networks will generally converge to a steady state, meaning that the range of their estimate on what the state in the next period will be has equal covariance. Accordingly, Shahrampour, Rakhlin and Jadbabaie (2013) model the state as a geometric random walk and study a wider variety of network architectures. Agents aim to minimize the quadratic loss function and measuring efficiency is done by estimating the variance between agents' estimate and the true state. They find a similar balance between the size of the network and estimating precision. On the one hand, more connections means an increase in prediction power. On the other hand, most efficient social learning is done by highly independent agents (with least common neighbors). Accordingly, using a Bayesian approach, I find that the adaptation mechanism operates most efficient if connections are limited.

[6] Chandrasekhar, Larreguy and Xandri (2015) empirically establish whether a group of people act more according to DeGroot learning or Bayesian models. Both types of learning are able to adequatly describe the data, but DeGroot learning generally is a better fit.

3 The Model

There is a finite set of agents, $i \in N = \{1, ..., n\}$, that receive an investment opportunity in each period, $t \in \{0, 1, ..., T\}$. In every period t, all of the agents have to decide for themselves and simultaneously on an action, $x_{it} \in X = \{0, 1\}$. The fixed set of agents that start the game continue playing until the game ends. The objective of each agent is to maximize individual VNM-utility $u_{it}(\theta_t, x_{it}) = (\theta_t - c)\, x_{it}$. If the investment is rejected, denoted $x_{it} = 0$, the payoff is zero, $u_{it}(\theta_t, 0) = 0$. If the agent decides to invest she has to incur cost, $c = \frac{1}{2}$, and the payoff depends on the unknown state of the world, $\theta_t \in \Theta = \{0, 1\}$ (the bad and good state). Investing, denoted $x_{it} = 1$, is optimal in the good state, $u_{it}(1, 1) = \frac{1}{2}$. If the state is bad the investment is not able to cover the cost and the yield is negative, $u_{it}(0, 1) = -\frac{1}{2}$. At the end of the game the payoff for all periods is summed, all agent see whether their choices were correct, and they receive their payoff.

At the start of each period, nature determines the state of the world. Furthermore, the random variable $\varepsilon = \Pr(\theta_t \neq \theta_{t-1})$, with $\varepsilon \in [0, 1]$, determines whether the state is different than in the preceding period[7]. For $\varepsilon = 0$ the state is fixed, denoted θ_f. For $\varepsilon > 0$ the state is dynamic, denoted θ_t. Moreover, the probability that the state of the world changes is common knowledge. Therefore, agents always take into account the *threat* of a time-varying state. To illustrate this threat, a probability $\varepsilon = \frac{1}{6}$ means that at the start of each period there is a one-to-six chance that the state is different than in the preceding period.

All the information known to the agent up to time t is summarized in a belief, $\mu_{it} \in \Delta(\theta) = (0, 1)$. The belief is the probability that the true state is good based on the history, $\mu_{it} = \Pr(\theta_t = 1 | h_{it})$. At period $t = 0$ there is no information, both states of the world are equally likely, and the initial belief is a flat prior for all agents, $\mu_{it_0} = \frac{1}{2}$. The agent is myopic as her memory can only contain a single value that carries over from one

[7] In essence, this is a Markov transition matrix with identical diagonal values.

period to the next. After the decision is made, the gathered information is discarded and the belief parameter reflects everything that is known up to that time. The investment decision maps the posterior belief over both states of the world into an action, $\sigma_{it} : \Delta(\theta) \rightarrow X$. The optimal one-period strategy, σ^*_{it}, is the choice for the action that maximizes expected payoff:

$$\sigma^*_{it}(\mu_{it}) \begin{cases} \mu_{it} > \frac{1}{2} & \rightarrow \quad x^*_{it} = 1 \\ \mu_{it} \leq \frac{1}{2} & \rightarrow \quad x^*_{it} = 0. \end{cases}$$

Over time, the individual strategy profile is a sequence of strategies, $\sigma^*_i = \{\sigma^*_{it}\}_{t \in T}$, that determine the optimal action in each period given the realized history. Moreover, the strategy profile, $\sigma^* = (\sigma^*_1, ..., \sigma^*_n)$, determines all actions that are taken by all players in every period (Tadelis, 2013).

To update her belief over the state of world the agent can consult two sources of information that form the history, $h_{it} = \{s_{it}, Z_{it}\}$, where $h_{t_0} = \{\emptyset\}$. First, starting from $t = 0$, each agent receives a private signal, $s_{it} \in S = \{0, 1\}$, that is conditional on the state of the world and independently and identically distributed. Moreover, the probability of receiving a correct signal is equal for all agents so that $q = \Pr(s_{it} = 1|\theta_t = 1) = \Pr(s_{it} = 0|\theta_t = 0)$, and the signal is noisy but informative, $q \in \left(\frac{1}{2}, 1\right)$. The probability of receiving an incorrect signal is $1 - q = \Pr(s_{it} = 1|\theta_t = 0) = \Pr(s_{it} = 0|\theta_t = 1)$. The signal could be interpreted as a report on the projected profitability of the investment, send to every potential investor. Each recipient has her own interpretation of this information or might have concerns about the trustworthiness of the sender. Hence, the report gives a first indication of the true state but cannot eliminate all uncertainty.

As a second piece of evidence, starting from $t = 1$, the agent observes the actions taken by neighbors. The extend to which her decision is influenced by others depends on the quality and quantity of this information. To assess the quality of the decision of others, the *reliability of neighbors* is the probability that the observed action of a neighbor, z_{jt}, is equal to the true state, $p = \Pr(z_{jt} = 1|\theta_t = 1) = \Pr(z_{jt} = 0|\theta_t = 0)$, where $p \in \left[\frac{1}{2}, 1\right]$.

This exogenous variable represents a predetermined level of trust in society and is assumed to be identical for all agents. Similar to Golub and Jackson (2010) and DeMarzo, Vayanos and Zwiebel (2003) the reliability is a subjective evaluation of the actions of others that does not reflect the true (unknown) percentage of correct observed actions. As an example, if $p = \frac{7}{10}$ and an agents sees her neighbor investing, the agent will always believe there is a 70 percent chance that this action is correct and reflects the true state at that time. Moreover, if $p = \frac{1}{2}$ observed actions are considered worthless and any information from the network is ignored, whereas for $p = 1$ everyone is convinced others always take the correct action, even if the actions of two neighbors are opposed (as these will cancel each other out). Valuing what others do enables social learning and the reliability of others is necessary to obtain a weighted average of the opinion of others. An increase in imposed reliability increases the influence of the actions of neighbors on the belief of the agent. Golub and Jackson (2010) justify this limitation of rationality, interpreting it as social pressure or willingness to match the actions of others. Furthermore, feeding the system with (partly) incorrect information does not have to be unrealistic as learning in the real world is often done by erroneous and dubious information. In Section 5.2 it is shown that, concerning the initial convergence of actions, agent in a complete network benefit from the combination of a flat prior and this imposed reliability, compared to a setting without neighbors. Moreover, the implications of these assumptions are discussed in Section 7.

The quantity of observed actions is reflected in the observation set, Z_{it}, a subset of all actions taken in the preceding period[8]. The observation set is a personal decentralized source of information as the elements included in the set depend on the location of the agent in the network. An individual with more members in her neighborhood will have a larger set

[8]The set X describes the two actions available to the agent. Z_t is the set of realised actions over all agents that can be observed at time t. Z_{it} is the subset of these realised actions visible to agent i at time t. I use this notation to dinstinghuish between the action the agent takes ($x_{it} \in X$) and what she observes ($z_{jt} \in Z$ for all $j \in N_i$).

of observations available and two agents have identical sets if they observe the same people. The observation set can be partitioned, $Z_{it} = \{Z_{it}^1, Z_{it}^0\}$, such that Z_{it}^1 contains the observed investments ($z_{jt} = 1$) and Z_{it}^0 the observed rejections ($z_{jt} = 0$). The difference between the number of elements in these sets is the *peer group dominant action*, $\delta_{it} = |Z_{it}^1| - |Z_{it}^0|$. This integer captures how many more neighbors have chosen to invest, and is negative if abstaining is the dominant choice among neighbors[9].

The interplay between the number of (Bernoulli distributed) observed actions and their weight (p) provides evidence on the true state of the world. The probability of a observing a certain number of investments conditional on the good state is Binomially distributed (for $|N_i| \geq 2$):

$$\Pr\left(\left|Z_{it}^1\right| |\theta_t = 1\right) = \binom{|N_i|}{|Z_{it}^1|} p^{\left|Z_{it}^1\right|} (1-p)^{\left|Z_{it}^0\right|}.$$

Accordingly, the probability of observing the same number of investments in the bad state is:

$$\Pr\left(\left|Z_{it}^1\right| |\theta_t = 0\right) = \binom{|N_i|}{|Z_{it}^1|} (1-p)^{\left|Z_{it}^1\right|} p^{\left|Z_{it}^0\right|}.$$

Before actually observing the actions of neighbors the agent can determine how many investments she should observe so observations will not influence her belief: disregarding the signal, the prior is equal to the posterior. She has an expectation of the size of δ_{it} and observing more (less) people invest will increase (decrease) her belief.

An important precondition for finding detrimental herd behavior is that every agent is boundedly rational. The actions taken by others are observed, but the agent is unable to infer why others take this decision and what this says about their information. Moreover, agents know their neigh-

[9]The possibility of imitation is overlooked by the agent. Mimicking is one of the consequences of this learning setup, but agents are unaware that their neighbors might not make an independent judgement. This simplifies the approach taken by Orléan (1995) in which the agent puts a subjective weight on the peer group opinion by estimating the number of imitators and independent agents and discards actions that are likely to be the result of mimicking.

bors, but are not informed about the structure of the network and their exact location. Strategic behavior is therefore ruled out. Bala and Goyal (1998) justify this assumption by noting that observing direct neighbors is cheap, but making inferences about their actions is infinitely expensive. If agents would take into account the network architecture they would have to be able to do complex calculations (see Gale and Kariv (2003) for fully rational individuals in a small network). DeMarzo et al. (2003) describe how an agent would have to distinguish old (correlated) from new information and must recall information received from neighbors in preceding periods. In the present paper the agent has no memory of previous signals and observations, other than the information aggregated in the belief parameter. In addition, the agent is myopic. She is unaware that the signal (from period t) and the observed actions (from period $t-1$) are generated in different periods. The observations are received in period t and therefore treated as evidence on the true state at that time.

Local interaction enables individuals to gather additional information in order to confirm or refute their findings from the report. The static network is represented by a graph with nodes and connections (N, g). A connection (or link) between i and j is denoted $g_{ij} = 1$. The present paper is limited to undirected networks, where interaction is reciprocal and a link g_{ij} automatically implies $g_{ij} = g_{ji}$. Furthermore, the neighborhood of agent i, denoted $N_i(g) \subseteq \{1, \ldots, i-1, i+1, \ldots, n\}$, consists of all the people she observes. Two network properties are worth mentioning[10]. First, the network is connected. Every node has at least one connection and there is a path from each agent to every other agent. Connected networks are analyzed because groups of unconnected components can be analyzed as a collection of connected networks (Bala and Goyal, 1998). Second, the agent's own action does not belong to the neighborhood, $i \notin N_i(g)$, and the extended neighborhood (neighbors of neighbors) is not observed. Interaction is restricted to those neighbors directly linked to the agent.

[10] For a detailed discussion of network properties see Goyal (2007) and Jackson (2008).

It is now possible to show the (in)efficiencies of this social learning setup.

4 The Evolution of Belief

In this section, I introduce the main building blocks that guide the analysis. The actions chosen by the agent, $x_{it} \in X$, are derived from her belief, $\mu_{it} \in \Delta$. Here, the main focus is on the values these beliefs can take on, the subset of beliefs that can be ruled out in advance and the way beliefs become active after an individual has decided to ignore information. First, there is a strict upper and lower limit on belief due to the continuous threat of an unobservable state change (ε). In addition, if beliefs cross a certain threshold, agents attain a level of certainty which will urge them to ignore new information and stop the belief from increasing or decreasing further (an information cascade). It is shown that the size of the neighborhood ($|N_i|$), the reliability of neighbors (p) and the signal precision (q) determine the value of this threshold. Furthermore, this disregard for new information will last for a limited number of periods, depending on the volatility of the state (cascades are temporary). After some time agents look for new information, thereby opening up the possibility of adaptation to the opposite state.

4.1 The likelihood ratio

The log likelihood ratio (henceforth: LLR) is the ratio of belief in the good state relative to the bad state expressed in logs. The LLR is comprised of three variables that form the posterior belief: the prior ratio, the signal precision ratio and the network confidence ratio[11]. The prior is the ratio

[11] $\lambda_{it+1} = \zeta_{it} + \chi_{it} + \lambda_{it} = \ln \frac{P\left(\theta_t=1|s_{it},Z_{it}^1\right)}{P\left(\theta_t=0|s_{it},Z_{it}^1\right)} = \ln \frac{P(s_{it}|\theta_t=1)}{P(s_{it}|\theta_t=0)} + \ln \frac{P\left(Z_{it}^1|\theta_t=1\right)}{P\left(Z_{it}^1|\theta_t=0\right)} + \ln \frac{P(\theta_t=1)}{P(\theta_t=0)}$.

of beliefs that carries over from the preceding period. This gives

$$\lambda_{it} = \ln\left(\frac{\mu_{it}}{1-\mu_{it}}\right), \tag{1}$$

where at the start of the game both states are equally likely, implying $\lambda_{it_0} = 0$ for all agents. Second, the signal precision ratio is given by

$$\begin{aligned} \zeta_{it} &= \alpha\left(2s_{it}-1\right), \qquad (2)\\ \text{where} \qquad \alpha &= \ln\left(\frac{q}{1-q}\right). \end{aligned}$$

The logarithm of the signal precision ratio is a constant value, positive if $s_{it} = 1$ and negative if $s_{it} = 0$. Third, the agent will take into consideration whether the actions of others are in accordance with her personal belief. The influence of neighbors is determined by the reliability of others (p) and the preferred action in the neighborhood (δ_{it}). The network confidence ratio is therefore given by

$$\begin{aligned} \chi_{it} &= \beta^{\delta_{it}}, \qquad (3)\\ \text{where} \qquad \beta &= \ln\left(\frac{p}{1-p}\right)\\ \text{and} \qquad \delta_{it} &= \left|Z_{it}^{1}\right| - \left|Z_{it}^{0}\right|. \end{aligned}$$

For a given number of neighbors in a static network this network confidence ratio is bounded.

The posterior likelihood is the sum of (1), (2) and (3). The posterior is the basis for the choice of action for an individual agent and includes all the information available at that time:

$$\lambda_{it+1} = \alpha\left(2s_{it}-1\right) + \beta^{\delta_{it}} + \lambda_{it}. \tag{4}$$

The balance between signal precision (q, included in α) and the reliability of neighbors (p, contained in β) is the main invariant measure that affects the equilibrium outcome. These variables measure the precision of the two sources of information. If the actions of others are not very reliable, the information included in a report is regarded as the most important source of information and a signal opposing the belief can alter the action. Nevertheless, a trustworthy report can be overruled if a sufficient number of agents agree on one of the two states. As the peer group dominant action (δ_{it}) increases, the neighborhood will become the primary source of information.

The belief of the agent can be derived from the LLR[12]. For a fixed state of the world ($\varepsilon = 0$) the belief is derived as follows

$$\mu_{it}(\theta_f) = \frac{e^{\lambda_{it+1}}}{1 + e^{\lambda_{it+1}}}. \tag{5}$$

If the state of the world is fickle ($\varepsilon > 0$) the prior is first updated according to (4) and then amended to take into account the possibility of a dynamic state of the world. After Bayesian updating to the posterior likelihood, the belief is

$$\mu_{it+1}(\theta_t) = (1-\varepsilon)\left(\frac{e^{\lambda_{it+1}}}{1 + e^{\lambda_{it+1}}}\right) + \varepsilon\left(1 - \frac{e^{\lambda_{it+1}}}{1 + e^{\lambda_{it+1}}}\right). \tag{6}$$

The threat of an underlying state change make the agent discount the information from earlier periods and introduces a continuous demand for new information. Information depreciates as a signal received several periods before might be conditional on a state of the world that is no longer up-to-date (Moscarini et al., 1998). From (6) it is easy to confirm that the threat of a change in state leads to a strict upper and lower bound on attainable beliefs. The dynamic belief is bounded $\mu_{it}(\theta_t) \in [\varepsilon, 1-\varepsilon]$ for $\varepsilon > 0$. Moreover, the findings by Moscarini et al. (1998) for high values of ε can be replicated. For $\varepsilon > \frac{1}{2}$ beliefs oscillate and actions alternate between zero and one. For $\varepsilon = \frac{1}{2}$ all information is superfluous which means $\mu_{it} = \frac{1}{2}$ for

[12] Smith and Sørensen (2000) use the term 'likelihood analogues' to emphasize that there is a direct one-to-one correspondence between the belief parameter and the LLR.

all agents in every period. These results deal with extreme values of ε and are omitted by the following assumption.

Assumption A *The volatility of the state is limited to $\varepsilon \in \left[0, \frac{1}{2}\right)$ from now on. There is a state of the world that changes, and it might change rapidly, but it is not likely to change two periods in succession.*

All variables in the LLR exhibit monotonicity. The signal is informative so that more good signals increase the likelihood of the good state. The signal precision ratio describes a relationship between two probability mass functions and satisfies the monotone likelihood ratio property (MLRP): $\frac{\partial}{\partial q}(q/(1-q)) = 1/(1-q)^2 > 0$ for $q \neq 1$. Moreover, the network confidence ratio is strictly increasing in the number of observed investments: $\frac{\partial}{\partial \delta_{it}}\left((p/(1-p))^{\delta_{it}}\right) = (p/(1-p))^{\delta_{it}} \ln(p/(1-p)) > 0$ for $\frac{1}{2} < p < 1$, and monotonically increasing in p: for a given δ_{it}, $\frac{\partial}{\partial p}\left(p^{\delta_{it}}/(1-p)^{\delta_{it}}\right) > 0$ for $p \neq 1$. Finally, belief is a continuous variable and the prior likelihood is a ratio of two probability density functions, $f(\mu|\theta = 1)$ and $f(\mu|\theta = 0)$, with corresponding cumulative density functions (CDF) $F^{\theta=1}(\mu)$ and $F^{\theta=0}(\mu)$, where for any $\mu_{it} \in (0,1)$, $F^{\theta=1}(\mu) \leq F^{\theta=0}(\mu)$ (first-order stochastic dominance) (Chamley, 2004). The agent will at least receive the utility from the state being bad, $\Pr\left(u_{it} \geq -\frac{1}{2}\right) = 1$, and there is a strictly smaller probability that she will receive more than that, $\Pr(u_{it} \geq 0) < 1$.

4.2 Herding and cascades

The process of social learning can give rise to detrimental network externalities if individuals ignore private information and, as a result, everyone else receives less information. An individual takes the action that is optimal from her personal point of view, but for an outsider the behavior may seem irrational due to the resulting loss of information for society. Therefore, it is important to define what this loss of information looks like, i.e. the difference between collective herding and individual cascades[13].

A *herd* is a characteristic of the actions taken by the population of agents (a choice dynamic). During a sequence of periods all actions for all agents are identical. There is conformity over which of the two states is likely to be true but agents can still adjust their belief. Therefore, they might choose to oppose the herd if their information is convincing and herding stops when the first agent decides to deviate. Moreover, a herd that does not stop, such that everyone agrees until the end of the game, is defined as a persistent herd.

Definition 4.1 *A herd is said to take place if, during a sequence $T_k \subset T$, where $T_k = \{t_{k_0}, ..., t_{k_u}\}$, $t_{k_0} \geq t_0$ and $t_{k_u} \leq T$, all actions are identical: for all $t_k \in T_k$, $x_{it} = x_{jt} = \bar{x}$ for all $i \in N$.*

Definition 4.2 *A herd is said to be persistent if actions are identical until the end of the game: $T_k = \{t_{k_0}, ..., T\}$, s.t. for all $t_k \in T_k$, $x_{it} = x_{jt} = \bar{x}$ for all $i \in N$.*

A herd describes the convergence of actions. For asymptotic learning beliefs have to converge as well, which is neither a result in the present paper, nor the goal of individuals in the present framework. A change of action to the true time-varying state demands a belief that is capable of change.

Beliefs are steadfast during an *information cascade* which, in contrast to the collective nature of herds, describes a mimicking process for individual agents (a belief dynamic). Cascades are detrimental because adaptation to the opposite state is not possible. During a cascade the opposite state is so unlikely that people decide on their action without the use of infor-

[13] As there is some ambiguity about what herds and cascades are, the starting point is Bikhchandani, Hirshleifer and Welch (2008) and Çelen and Kariv (2003). These definitions, related to one-time sequential decision making, are applied to repeated simultaneous decisions. Moreover, these definitions are not in accordance with other literature. As an example, Chamley (2004:64) defines the start of an information cascade as the moment when the agent does not follow her signal but still processes the information and the start of a herd as the moment when agents ignore private signals.

mation, as it is regarded as redundant. The individual might save time by not reading the report and save effort by not having to meet with others to hear their decision. The learning process has ceased as there is no Bayesian updating of the LLR and the agent mimics the action taken in the preceding period.

The *cascade set* is the union of the upper and lower cascade set, a range of beliefs for which the ex ante probability of a different action than in the preceding period is zero, $\mu_{it} \in H_l \cup H_u \Rightarrow \Pr(x_{it} \neq x_{it-1}) = 0$. A cascade starts in period t when the posterior belief enters one of the cascade sets in period $t-1$. To find the cascade set, the belief is partitioned $\mu_{it} \in H_l \cup H_m \cup H_u = [0,1]$. The *lower cascade set*, $H_l = [0, \eta)$, is a set of prior beliefs in a specific round for which the agent will always conclude not investing is optimal, $x_{it} = x_{it-1} = 0$. Even if all information is opposing this belief, $Z_{it}^1 = \{N_i\}$ and $s_{it} = 1$, the range of feasible posteriors is $\mu_{it+1} \in \left[0, \frac{1}{2}\right]$. Similarly, for a belief in the *upper cascade set*, $H_u = (1-\eta, 1]$, the posterior belief will always indicate investing is optimal, $\mu_{it+1} \in \left(\frac{1}{2}, 1\right] \Rightarrow x_{it} = x_{it-1} = 1$, even if all information indicates otherwise, $Z_{it}^1 = \{\emptyset\}$ and $s_{it} = 0$. In between, $\mu_{it} \in H_m = [\eta, 1-\eta]$, beliefs are *active* as the agent chooses both actions with positive probability and uses her information to make this choice, $\Pr(x_{it} = 1) < 1$.

Definition 4.3 *During an information cascade the agent mimics her own action taken in the previous period,* $\Pr(x_{it} = x_{it-1}) = 1$. *The range of possible posterior beliefs is restricted to the same half of the interval as the prior belief. Any information is discarded.*

The threat of a change in the state of the world (Assumption A) and the cascade set (Definition 4.3) restrict the evolution of belief. Figure 1 gives a graphical representation of the cascade set and the range of unfeasible beliefs due to a volatile state and serves as a template for graphs in the next section.

Active beliefs are a necessary condition for agents to adapt. If cascades

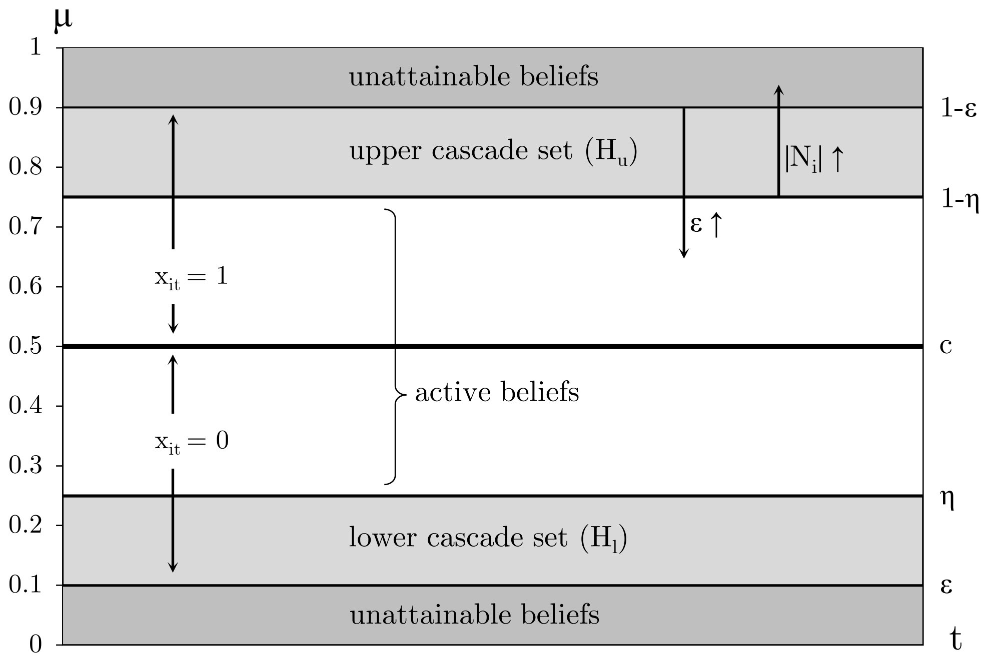

Figure 1. Boundary conditions for the evolution of belief over time. Beliefs are active if there is a need for new information and both actions are possible. In the cascade sets information is ignored, but beliefs bounce back towards the center after some time due to the depreciation of information. Extreme beliefs are unattainable as a result of the dynamic nature of the state of the world.

would be everlasting, beliefs would be constant until the end of the game, thereby preventing a change in action. Furthermore, it is easy to see that the cascade region is unattainable if the state is highly volatile ($\varepsilon \geq \eta$). Then, as a result of the high rate of depreciation, there is a constant demand for new information. Agents will use all the information they receive to make a balanced decision. Moreover, whether the belief of the agent is in a cascade set depends on the size of the neighborhood. For an increase in the number of neighbors ($|N_i|$) the cascade threshold (η) decreases. Ignoring the signal plus a group of others requires a stronger belief than ignoring just the private signal. Therefore, Lemma 1 shows that an agent with more neighbors is less susceptible to cascades.

Lemma 1 *The upper bound on the lower cascade set, $H_l = [0, \eta)$, satisfies $\eta = \left(1 + \frac{q}{1-q}\frac{p}{1-p}^{|N_i|}\right)^{-1}$. Hence, the size of the cascade set $H_l \cup H_u = [0, \eta) \cup (1-\eta, 1]$ is decreasing in $|N_i|$ for $|N_i| > 0$* and $\lim_{|N_i| \to \infty} \eta = 0$.

Smith and Sørensen (2000) show that a necessary condition for the existence of cascades is that beliefs are bounded. Beliefs are bounded if there is a maximum amount of information that can be derived from the signal (and in this case observations), which causes the LLR to be finite and bounded away from zero (Acemoglu et al., 2011). For an agent without neighbors with signal precision q, the signal precision ratio is $\zeta_{it} \in \{-\ln(q/(1-q)), \ln(q/(1-q))\}$. As $q \in \left(\frac{1}{2}, 1\right)$ this means $0 < \ln(q/(1-q)) < \infty$, which shows that the maximum distance between a prior and posterior belief in a specific round is finite and depends on the informativeness of the signal. In contrast, for an unbounded likelihood ratio signals have to be arbitrarily strong (a continuous signal space). As a result, the probability an agent chooses one of both states is positive for any $\mu_{it} \in (0, 1)$ as the information can always alter the action and the prior is overwhelmed by the signal. Hence, the only cascade sets are $\mu_{it} \in \{\emptyset\} \cup \{1\}$ and learning will continue until beliefs have converged to the true state (Acemoglu et al., 2011).

During an information cascade signals and observations are ignored. The Bayesian posterior belief (4) is then replaced by

$$\mu_{it+1}(\theta_t) = (1-\varepsilon)\mu_{it} + \varepsilon(1-\mu_{it}) \tag{7}$$
$$\text{for} \quad \mu_{it} \in [0, \eta) \cup (1-\eta, 1].$$

For a fixed state we have $\varepsilon = 0$ and thus $\mu_{it+1}(\theta_f) = \mu_{it}$. Therefore, a belief that is in a cascade set will be constant until the end of the game as the individual has collected information, made up her mind, and is certain that this is the optimal action. Changing the decision is unnecessary as there is no reason to suspect the outcome will be any different. The evolution of belief during a cascade subject to a dynamic state ($\varepsilon > 0$) is characterized by a discounted version of the prior (7), which evolves decreasing convex if $\mu_{it} \in (1-\eta, 1]$ or increasing concave if $\mu_{it} \in [0, \eta)$. Consider an individual with a belief that indicates she should invest $\left(\mu_{it} > \frac{1}{2} \Rightarrow x_{it} = 1\right)$. At some point in time she is isolated from the outside world and is unable to receive new information. Each period, some of the information she had collected becomes less valuable, as it is conditional on a state long ago. As time goes by, the information is reduced to a tiny clue about the current state. Her best option is to simply choose the same action she chose when she last received information, while uncertainty about the correctness of this decision increases. This illustrates the non-Bayesian movement of belief towards the fixed point, $\mu_{it} = \frac{1}{2}$, during cascades. The fixed point is the asymptote for any belief, $\mu_{it} \in [0, 1]$, inserted in (7) as $t \to \infty$ (the proof of Lemma 2 provides a detailed discussion on the fixed point).

In this framework ignoring information is a personal consideration such that others can still accumulate information. Agents can decide for themselves when they have a need for new information and when they discard signals and observations. Therefore, we can replicate the results of Moscarini et al. (1998) that cascades are temporary and will stop within a finite period of time. The behavior during a cascade is no different than for sequential decision makers, because the addition of a network structure has no influence on cascade behavior as observations are discarded. As soon as beliefs are active, $\mu_{it} \in [\eta, 1-\eta]$, the belief has exited the cascade region, signals and observations are regarded valuable again, and updating is again done by (4).

Lemma 2 *A cascade is temporary if the state of the world is dynamic and a network is present. Beliefs in the cascade set are adjusted in the direction of the fixed point, $\mu_{it} = \frac{1}{2}$, due to the depreciation of information. This process continues until the belief is active and the probability of both states of the world is positive or until the game has ended.*

The foregoing analysis has shown that (i) ignoring information is less likely when neighbors are present and (ii) ignoring information is restricted to a limited number of periods. Agent will have a tendency to ignore information as additional information is aggregated in the belief parameter but, as older information is less valuable, temporary cascades enable beliefs to become active after some time. No matter how convinced an agent is about the correctness of her choice, after some time there is a positive probability that both actions are chosen. In the next section I will use this finding to show that active beliefs enable agents to take up new information and change their action. The goal of the next section is to describe in what way beliefs, subject to the above conditions, evolve for a variety of network architectures. Hence, it is shown that cascades are no longer an obstacle to deviate from the herd, but interaction can prevent adaptation to a different state.

5 Main Results

The goal of the following analysis is to show in what way local interaction affects the ability of a group of people to react to actual changes in the state of the world. The belief is limited by the conditions laid down in the previous section, so that we can now observe what happens when a sequence of signals is generated for various network structures. The length of the game (T), the signal precision (q), the reliability of neighbors (p) and the probability of a state change (ε) are determined before the start of the game and fixed throughout. As a result, there are two factors that can cause differences in actions: the realized signals and the network architecture.

In contrast to the *threat* of a change in the state of the world (the probability ε is common knowledge) an *actual* change in the underlying state ($\theta_t \neq \theta_{t-1}$) is not directly visible for agents but will alter the distribution of signals. The expected value of signals in the good state is $E\left[s_{it}|\theta_t = 1\right] = q$. In the bad state we expect more negative signals $E\left[s_{it}|\theta_t = 0\right] = 1 - q$. Thus, for a static network, the outcome is determined by the realized sequence of signals. To isolate cause (the network architecture) from effect (the actions and beliefs of a group of agents), the signals are kept constant. Therefore, a sample path is one unique sequence of signals that causes one outcome[14]. An outcome is the value of belief for each agent and for every period corresponding to this sequence of signals, $\{\mu_{it}\}_{i\in N,t\in T}$. Any statement about the influence of the network structure on differences in actions assumes an identical sample path. However, the following results are robust to changes in the signal distribution. Furthermore, the following examples are all generated by one unique sequence of signals.

Example A Assume the state of the world is $\theta_{T_A} = 1$ for a sequence of periods, $T_A = \{t_0, ..., t_{k-1}\}$, and then changes to $\theta_{T_B} = 0$ for the rest of the game, $T_B = \{t_k, ..., T\}$. The associated expected value for a sequence of independent signals for all agents is q for $t \in T_A$ and $1 - q$ for $t \in T_B$.

Each following subsection discusses a network architecture. Individuals without neighbors (Section 5.1) use the signal as a sole source of information. This section serves primarily as an illustration of temporary cascades as herds will emerge, but not as a consequence of interaction. People in a complete network (Section 5.2) have the highest possible level of interaction and are therefore prone to (persistent) herds. For agents in an incomplete network (Section 5.3) time and location determine their ability to change actions.

5.1 Isolated agents or unreliable neighbors

For an isolated agent the private signal is the only variable affecting belief. An isolated agent is either a person without neighbors, $N_i = \{\emptyset\}$, or someone that receives information from the action of others but judges them as unreliable, $p = \Pr(z_{jt} = \theta_t) = \frac{1}{2}$. As a result, potential observations are discarded and the network confidence ratio will not influence beliefs, $\chi_{it} = 0$ for all $i \in N$. The LLR is therefore given by $\lambda_{it+1} = \zeta_{it} + \lambda_{it}$.

For isolated agents, switching actions to adapt to a different state is always possible. After all, every individual receives a sequence of signals

[14] The collection of sample paths constitutes every possible outcome and the probability of a sample path is the product of the probabilities of the independent signals.

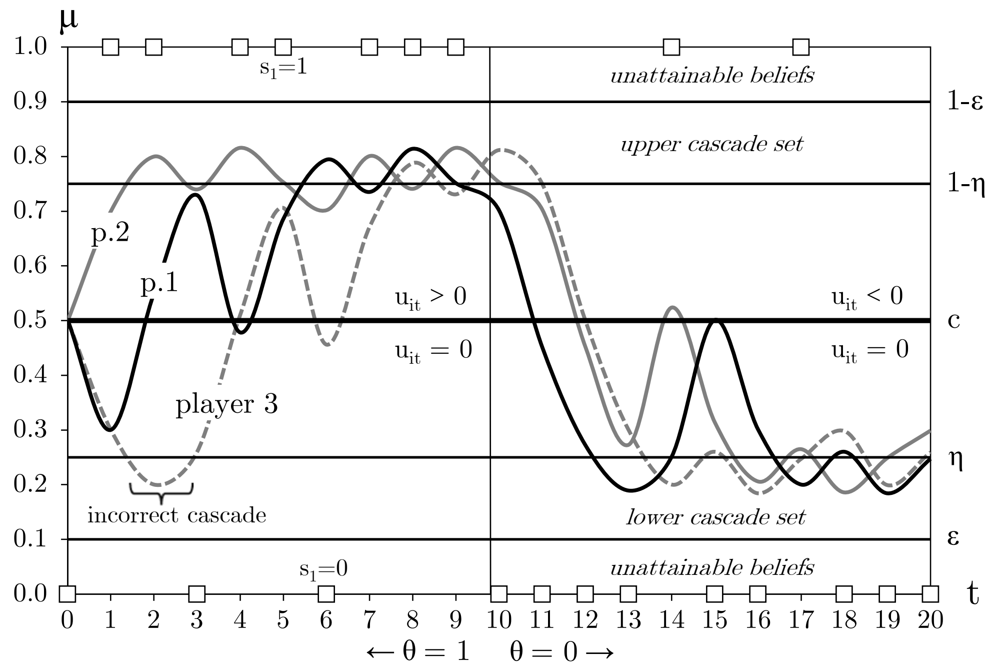

Figure 2. Three isolated agents successfully adapt to a change in the state of the world at $t = 10$ ($q = \frac{3}{4}$, $p = \frac{1}{2}$, $\varepsilon = \frac{1}{10}$). As an illustration of the distribution of signals, the white boxes represent the signals send to player 1. Signals for other agents are not shown. The fluctuations on the top and bottom of the graph are caused by temporary cascades.

that is informative and conditional on the actual state of the world as the only source of information. Moreover, the agent is aware that the state is dynamic, such that if initially the majority of information indicates that investing was optimal, and subsequently she receives messages that indicate investing can hurt payoffs, the agent will notice this change in tone and act accordingly. There is no reason to have doubts about this information, other than the usual noise (q).

Example A.1 The evolution of belief for three isolated agents ($N = 3$ and $p = \frac{1}{2}$), receiving a randomly generated sequence of signals, is shown in Figure 2. This is one of (2^{63}) sample paths for $q = \frac{3}{4}$ and $\varepsilon = \frac{1}{10}$. Furthermore, the actual state change from good to bad occurs at $t_k = 10$, which means investing is optimal for the first ten periods and rejecting the investment is optimal until $T = 20$. Moreover, the cascade threshold is not influenced by neighbors ($\eta_0 = 1 - q$) and the state is not very volatile

$(\varepsilon < \eta_0)$ so that the agents can achieve a level of certainty and beliefs can enter the cascade set.

Beliefs follow the change in underlying state. At the start of the game the LLR tends to ∞ and the belief tends to the upper bound, $1-\varepsilon$, for all agents. The expected value of the signal precision ratio is positive, $E\left[\zeta_{iT_A}|\theta_{T_A}=1\right] > 0$, for any $q \in \left(\frac{1}{2},1\right)$. After the state has changed, there is a sudden increase in negative signals, $E\left[\zeta_{iT_B}|\theta_{T_B}=0\right] < 0$, so that the LLR tends to $-\infty$ and the expected belief tends to ε for all agents.

There is a positive probability that agents take incorrect decisions and beliefs end up in an incorrect cascade. Thus, the payoff for investments will be negative if the agent decides to invests and state is bad, or there are foregone profits if not investing is chosen in the good state. The signals are informative, but there is a chance that an individual receives a series of bad reports or wrongly interprets the information for a number of successive periods. The probability of an incorrect cascade increases as the signal precision (q) decreases.

Example A.2 Consider player 3, receiving three incorrect signals at the start of the game. Her belief enters the cascade region at $t=2$ $\left(\mu_{32}=\frac{2}{10}<\eta\right)$ and all her actions until that moment are incorrect. She repeatedly decides not to invest, but could have made a profit. The probability of this incorrect cascade is the probability of receiving three consecutive incorrect signals, $(1-q)^3 > 0$.

During such an incorrect cascade, if the state is fixed, the agent ignores any further information and is therefore unable to rectify this decision. A dynamic state of the world enables her to recover from incorrect actions, as the depreciation of information ensures the agent will have a need for new signals after a finite period of time (see Lemma 2). After some time the belief is active (both states of the world are again true with positive probability) which enables the agent to adequately react to the signals she receives. Hence, she is again expected to take the correct decision.

Example A.3 The belief is bounded by the threat of a change in the unobservable state, $\mu_{it} \in [\varepsilon, 1-\varepsilon] = \left[\frac{1}{10}, \frac{9}{10}\right]$, and will bounce back once it is in the cascade region: $\mu_{it} \in [0, \eta) \cup (1-\eta, 1] = \left[0, \frac{1}{4}\right) \cup \left(\frac{3}{4}, 1\right]$. For all agents beliefs become active after a maximum of two periods after an information cascade has started[15].

These examples illustrate the mechanism that enables adaptation of beliefs and actions to the true state, irrespective of whether neighbors are present. Agents know the state may have changed, which makes the value of information depreciate. A person that ignores signals will have a brief look at some signals every now and then to see whether the action is still up-to-date. She will then change her action if she notices the changes in the signal distribution. In doing so, a sequence of signals that is conflicting with her belief causes her to reconsider and might lead to a different action.

A group of agents that has successfully adapted to a new state is in a herd (they all take the same action $\bar{x}$). Isolated agents do not receive meaningful information from others such that herding does not arise as a result of mimicking, but because they are all expected to successfully adapt to the true state over time. In contrast, a persistent herd, where everyone takes the same action until the end of the game ($T_k = \{t_{k_0}, ..., T\}$), is unlikely. Isolated agents freely adapt to the true state and after each change in the time-varying state one agent is the first to deviate, thereby breaking the herd. Provided that the underlying state changes, the chance of a persistent herd approaches 0 if $t \to \infty$. Proposition 1 describes this behavior.

Proposition 1 *An isolated agent, $N_i = \{\emptyset\}$, or agent with unreliable neighbors, $p = \frac{1}{2}$, will be able to adapt to a real change in the state of the world ($\theta_t \neq \theta_{t-1}$). The probability of a persistent herd approaches 0 if $t \to \infty$ and $\varepsilon > 0$.*

[15]There is a second strict bound that is not defined and is omitted from the rest of the paper. For a minimum (maximum) active prior belief $\mu_{it} = \eta = \frac{1}{4}$ ($\mu_{it} = 1 - \eta = \frac{3}{4}$) there is a maximum decrease (increase) in belief caused by the signal $s_{it} = 0$ ($s_{it} = 1$) which results in the posterior belief $\mu_{it} = 0.18 > \varepsilon$ ($\mu_{it} = 0.82 < 1 - \varepsilon$). This belief will exit the cascade region in two periods.

As a result of the absence of local interaction, agents decide individually and are not bound by instructions. They are expected to (generally) take the same actions and herding is likely, although any agent is free to deviate from the herd and there is no fixed order in which adaptation takes place.

5.2 Complete undirected networks

Appreciating the opinion of others will influence the decision making process. In a complete network everyone is connected with everyone ($|N_i| = n - 1$ for all agents). Moreover, the imposed level of trust in society is such that agents believe the majority of the decisions of their neighbors are correct, $p > \frac{1}{2}$. As the agents now have two sources of information, they have to weigh the precision of the signal (q) and the reliability of neighbors (p). Moreover, even for untrustworthy neighbors, the network takes over as primary source of information if the difference between the number of people investing and rejecting the investment is large. In a given period, the belief of the agent follows the direction dictated by the actions of others if:

$$
\begin{aligned}
& |\delta_{it}| > \frac{\alpha}{\beta}, \qquad (8)\\
\text{with} \quad & \alpha = \ln \frac{q}{1-q} \\
\text{and} \quad & \beta = \ln \frac{p}{1-p}.
\end{aligned}
$$

The peer group dominant action (the integer δ_{it}) varies over time as agents observe neighbors switching from investing to rejecting, and vice versa. The boundary α/β defines the number of agents that have to agree on a certain action before the signal is overruled by the peer group opinion, in the LLR $|\zeta_{it}| < |\chi_{it}|$. For agents agreeing with the majority, if (8) holds, posterior beliefs are constraint to the interval the prior is in, either $\mu_{it} \in \left[\varepsilon, \frac{1}{2}\right)$ or $\mu_{it} \in \left(\frac{1}{2}, 1-\varepsilon\right]$. A belief that causes a switch in actions is unfeasible. For those not agreeing with the dominant action, the belief will be drawn

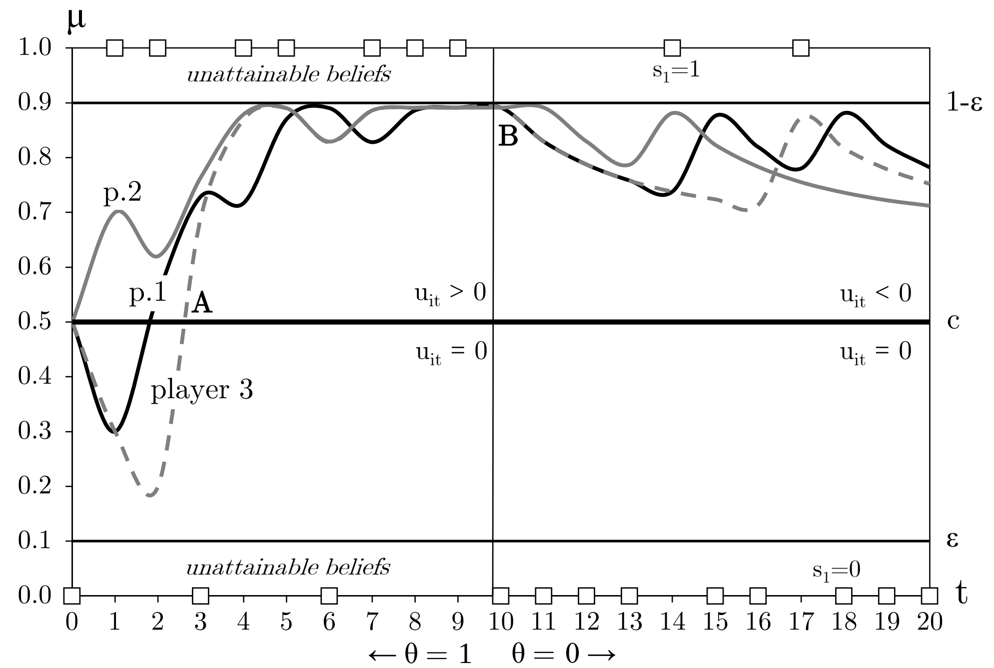

Figure 3. The evolution of belief for three agents in a complete network ($q = \frac{3}{4}$, $p = \frac{2}{3}$, $\varepsilon = \frac{1}{10}$). A persistent herd starts at A. Therefore, these agents are not able to adapt to the state change at B. Moreover, cascades are not possible due to the volatility of the state and the number of neighbors for each agent ($\varepsilon > \eta$). Therefore the cascade sets are not displayed.

towards the opposite side. No combination of signals is able to prevent their belief from conforming with the peer group opinion eventually. The only possibility to postpone a herd is when the belief of an agent is in the opposite cascade set, in which any information is bluntly ignored. Even so, ignoring information is limited to a finite number of periods and temporary cascades force the agent to join the others eventually.

To illustrate the behavior of agents in a complete network, Figure 3 displays the evolution of belief for the same agents, receiving the same sequence of signals, as in Examples A.1 to A.3. Here, the difference is that agents can learn from the actions of the people they observe.

Example A.4 In Figure 3 the network takes over as the primary source of information if the agent observes a difference of at least two neighbors taking the same action, $|\delta_{it}| \geq 2$. Notice that player 2 believes investing is optimal throughout the game. Moreover, at the beginning of the game player 1 and 3 are not influenced by neighbors. They both observe one neighbor investing and the other rejecting the investment, $|Z_{it}^1| = |Z_{it}^0| = 1 \Rightarrow \delta_{it} = 0$, and follow their signal leading them to an incorrect decision. Furthermore, after the first signal, player 1 correctly interprets all of the following signals, decides to switch actions, and believes that investing is optimal at period $t = 2$. From this moment on, player 3 observes two investments, $\delta_{3t} = 2$. Therefore, the actions of neighbors are considered to be a more important reflection of the true state than the signal and the belief is drawn towards the correct action in the following periods. As long as player 3 opposes the actions taken by the other two, a change in action is still possible for player 1 and 2. For them the signal is still the primary source of information, $\delta_{2t} = 1 \Rightarrow \zeta_{it} > \chi_{it}$. Thus, if one of these agents receives a series of (incorrect) reports, indicating that they should withdraw their investment, they are able to adjust their belief accordingly and avoid the herd that now starts at A.

If all agents agree on the same action and they are unable to deviate, a herd automatically implies a persistent herd. Formally,

$$|N_i| > \frac{\alpha}{\beta} \qquad \text{for all } i, \tag{9}$$

represents the threshold for the size of the neighborhood such that none of the agents is able to switch actions once action have converged. A neighborhood of this size gives rise to the *gridlock effect*: for all agents the reliability of the report (q) is not able to counteract the combination of the reliability of neighbors (p) and the peer group dominant action, $|\delta_{it}| = n-1$. There are simply too many neighbors agreeing on the same action. Agents will notice changes in the signal distribution but are unwilling to challenge the conventional wisdom and adapt to a changing state.

Example A.5 The persistent herd, that prevents adaptation to the true state in Figure 3, is caused by the gridlock effect. At A every individual observes two neighbors investing ($\delta_{it} = 2$ for all agents). As a consequence, everyone considers the network as the primary source of information and any report will be overruled by the (less reliable) observations. At $t = 10$ the state of the world changes and refraining from investing is optimal (B). Nevertheless, all agents continue to invest even though most of the reports they receive indicate otherwise.

Moreover, because of the uncertainty about the state of the world and the presence of two neighbors, beliefs will never enter the cascade region in this example ($\varepsilon > \eta_2$). As a result, cascades are nonexistent and every period the LLR will be updated by taking into account information from both sources.

To assess the impact of a persistent herd there is an important difference between short-term and long-term outcomes. Compared to isolated agents, a complete network increases the probability that everyone chooses the correct action in the short-term. In the first period everyone receives an informative signal without observing others. Similar to a fixed world setting, these first signals are important to determine the herd equilibrium for the rest of the game (see for example Gale and Kariv, 2003). The probability that initial herds are incorrect remains, but decreases as signal precision or the size of the complete network increases. Moreover, the assumption of a flat prior (all agents believe $\mu_{it_0} = \frac{1}{2}$) increases the likelihood that the correct decision is taken by most agents in the first periods. Not only do we expect most of them to take the correct action after the first signal, $E\left[\left|Z^1_{t_1}\right| |\theta_{t_1} = 1\right] = nq$ and $E\left[\left|Z^0_{t_1}\right| |\theta_{t_1} = 1\right] = n\left(1-q\right)$, this also increases the expected number of observed investments in period one: $E\left[\delta_{it_1}|\theta_{t_1} = 1\right] > 0 \Rightarrow E\left[\chi_{it_1}|\theta_{t_1} = 1\right] > 0$. Therefore, in addition to the signal, the peer group opinion directs initial beliefs towards the correct short-term action.

The drawback of local interaction is that long-term behavior is characterized by suboptimal actions. Agents will never change their action, even though they are aware that the state of the world (and thus their preferred decision) changes with positive probability. If we compare individual behavior during persistent herds and information cascades, a cascade occurs when there is a level of certainty over the state of the world that makes the agent ignore any further information. She knows that out of all the possible combinations of signals and observations there is no combination that is going to change her belief on the optimal action. In contrast, during a persistent herd, the agent receives lots of information that can potentially cause a change in action. The problem is that everyone is waiting for their neighbors to change and, as a consequence, none of the agents is willing to switch actions. There is a *potential* combination of signals and content of the observation set that can alter the action. The gridlock effect occurs because none of the neighbors will *actually* take the actions that will lead to these observations.

These results can be extended to a complete network with any number of nodes without loss of generality. Proposition 2 describes the persistent herd behavior in complete networks.

Proposition 2 *In a complete network, if* $|N_i| > \dfrac{\alpha}{\beta}$ *and* $\varepsilon > 0$, *actions converge. Moreover, the probability of a persistent herd approaches* 1 *if* $t \to \infty$. *If* $|N_i| \leq \dfrac{\alpha}{\beta}$ *there is a positive probability that any herd is temporary and agents are able to adapt to changing conditions with positive probability.*

A persistent herd is a manifestation of the curse of information. There is too much information in the economic system. Therefore, people are unable to distinguish the signal that contains valuable information on the true state and there is too much weight on the actions of peers. The characteristics of a persistent herd and the finding that individuals overweigh their information are closely related to the concept of *correlation neglect*: individuals underestimate the level of correlation of their various information sources, treating them as independent, and causing conformity in behavior (Levy and Razin, 2015). Aggregate group behavior results in outcomes that are different than if agents would only have private information. Moreover, DeMarzo et al. (2003) introduce the concept of *persuasion bias* to indicate that agents fail to account for possible repetitions of information they receive. Every piece of information is seen as new and influences the belief as such. Furthermore, if individuals threat all information as new, the level of connectedness of agents (their social influence) determines their influence on the outcome. During a persistent herd, no one is able to infer that the actions they take influence the actions they observe and that everyone receives information from the same pool of peers. Eyster and Weizsäcker (2011) perform several experiments and find that most people neglect correlation when deciding on the composition of their investment portfolio. If all investors overestimate their biased information this can lead to excessive risk taking in financial markets (see also Enke and Zimmerman (2013)).

5.3 Incomplete undirected (line) networks

Limiting the number of connections for some of the agents restricts the flow of information and can make people that would follow the herd in a complete network susceptible to their signal again. In this section I give an example of an incomplete network architecture and show that this enables the group to process new information and adapt to changes. The starting point is that every incomplete network is a subset of the complete network. In an incomplete network some of the links are removed and at least one of the sets in the observation set $Z_{it} = \{Z_{it}^1, Z_{it}^0\}$ contains less elements.

As a special case of the incomplete network, a *line network* consists of two groups of nodes (2 outer nodes and $n-2$ inner nodes) that differ in the number of connections. Nodes at the ends of the line have just one link and one neighbor. In between, all agents have two neighbors. As a result, the observable peer group dominant action is limited to integers in the interval $\delta_{it} \in [-2, 2]$. Moreover, individuals with one connection receive less information than agents with two connections. Therefore, the size of the cascade region is larger for the outer nodes $\left(\eta_1 > \eta_2\right)$. Based on the

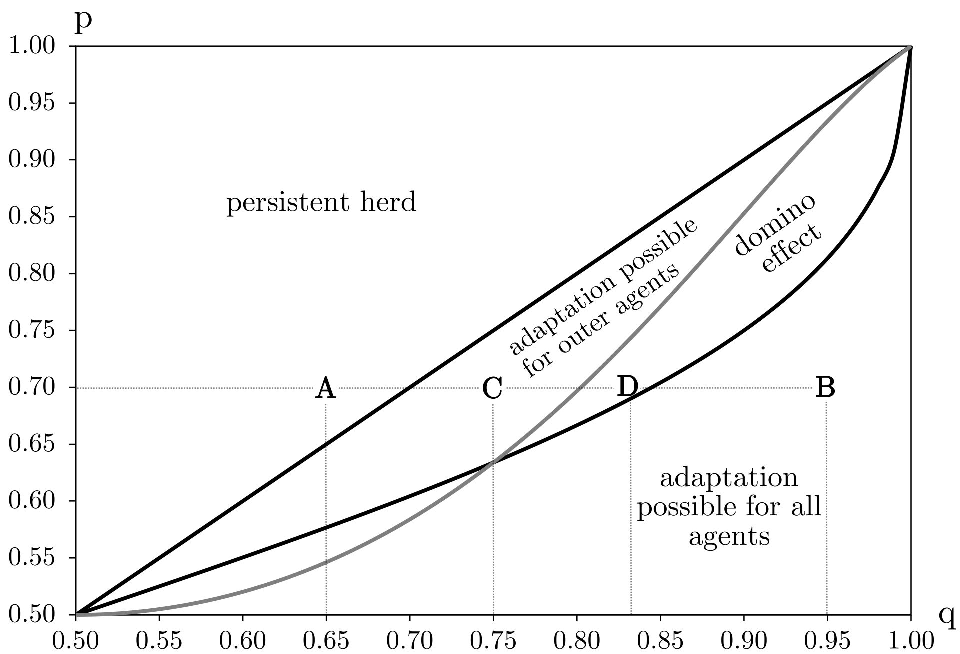

Figure 4. Differences in behavior for a three-person line network depending on signal precision (q) and the reliability of neighbors (p). In region A none of the agents is able to adapt. In region B everyone can adapt. In region C agents with one connection are able to adapt but are not expected to do so because of their signal. In region D the signal is strong enough and we expect a domino effect after each change in the underlying state of the world.

ratio of signal precision (q included in α) and reliability of neighbors (p included in β) we can distinguish three intervals illustrated in Figure 4.

If one observation is enough to overrule the signal ($0 \leq \alpha \leq \beta$) any herd is persistent and adapting to changing circumstances is impossible once actions converge. Then, behavior is similar to the gridlock effect found in the complete network as $q \leq p$ automatically implies a persistent herd (9). Point A in Figure 4, at which 65 percent of reports are interpreted correctly, but 70 percent of the actions of any neighbor are judged to be correct, is one such point. Every agent will follow the action(s) of her neighbor(s). Furthermore, if observing more than two neighbors is necessary before the gridlock effect occurs ($\alpha > 2\beta$) the ability to adjust is limited. Beliefs are influenced by the observation of neighbors, and agents need convincing signals before they are willing to switch actions. Nonetheless, trying a different action is always possible for all agents, a gridlock is impossible and a persistent herd is unlikely. At point B the signal is very reliable and reflects the true state of the world 95 percent of the time. Even though 70 percent of the action of others are reliable and some agents observe two corresponding actions, the signal guides beliefs as the primary source of information for everyone.

The location of the agent in the network and the moment in time are important when players in the center are not able to switch actions, while the other two players can change. Then, the network dynamics determine whether and when someone is able to switch actions. This intermediate situation is illustrated at point C. Here, the signal (correct 75 percent of the time) can overrule one but not two corresponding observations (of which 70 percent is considered correct). As a result, the signal is the primary source of information for agents with one neighbor and the network determines what to believe for agents observing two similar actions. This dynamic occurs when $1 \leq \alpha \leq 2\beta$. This gives[16]

$$p \leq q \leq \left(1 + \frac{(1-p)^2}{p^2}\right)^{-1}. \tag{10}$$

Contrary to a complete network, there is no guarantee that actions converge when (10) is satisfied. Persistent herds are not the unique outcome as the two outer agents will always be able to react to signals. Nevertheless, they are drawn towards the herd by their neighbor. At C, on average, there are not enough signals and they are not precise enough to let an outer player escape the herd. If the true state of the world is $\theta_t = 0$ ($\theta_t = 1$) and everyone is taking the opposite action, the outer players are only expected to change when their LLR tends to $-\infty$ (∞), implying $E\left[\zeta_{it}|\theta_t = 0\right] + \beta < 0$ $\left(E\left[\zeta_{it}|\theta_t = 1\right] + \beta^{-1} > 0\right)$, and can oppose their only neighbor if signals are

[16] Accordingly $\left(1 + \sqrt{1/q - 1}\right)^{-1} \leq p \leq q$ for $q \in \left(\frac{1}{2}, 1\right)$ and $p \in \left(\frac{1}{2}, 1\right)$.

considered to be important enough:

$$p < \left(1 + \left(\frac{q}{1-q}\right)^{(1-2q)}\right)^{-1}. \tag{11}$$

If (10) holds, one or more agents are able to escape the herd as they consider their signal important while the others are unable to switch actions. If (11) holds, the expected sequence of signals (and their precision) generates enough evidence for these agents that the opposite action delivers a higher payoff. This situation is illustrated at D. On average, there are enough signals indicating that the outer players have to switch actions, contradicting the observations they receive from their neighbor. One of them is likely to change actions and subsequently the agent in the center observes $\delta_{it} = 0$, which means she is able to follow.

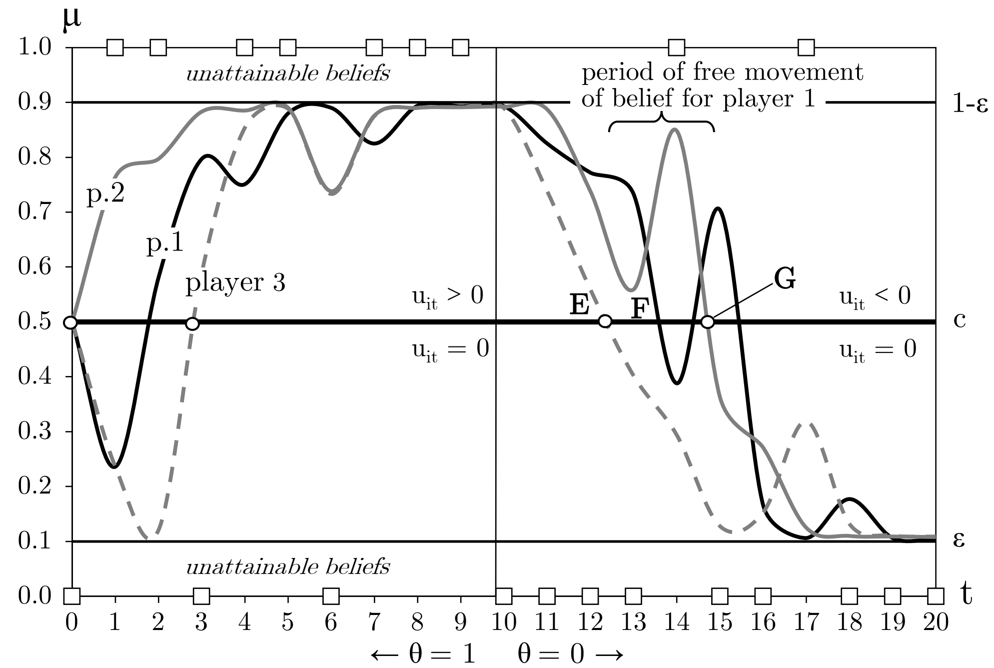

Figure 5. The domino effect for a three-person incomplete line network ($q = 0.83$, $p = \frac{7}{10}$, $\varepsilon = \frac{1}{10}$). One of the outer players (2 and 3) has to change actions before player 1 can follow. A first domino effect is visible from $t = 0$ and ends when the herd begins ($t = 3$). A second domino effect starts as the state changes ($t = 10$) and ends when player 1 definitively rejects the investment ($t = 16$).

Figure 5 displays an example of a three player line network, receiving the same signal sequence as in Examples A.1 to A.6. To ensure adaptation by the outer agents, the reliability of neighbors has been increased to $p = \frac{7}{10}$ and 83 percent of signals are correct (corresponding to D in Figure 4).

Example A.6 In Figure 5 actions have converged after ten periods so that investment is the preferred choice of action for all agents. Following this herd the underlying state changes, which requires adaptation to prevent a loss. Player 1 is in the center of the line network. She observes two investments ($\delta_{1t_{10}} = 2$) and is therefore unable to change actions. Player 2 and 3 observe the center player investing ($\delta_{it_{10}} = 1$). At point E player 3 has had several reliable reports that she should change her action to prevent a loss. She deviates to adapt to the correct state. After the deviation is observed, the content of the observation set of player 1 has changed ($\delta_{1t_{13}} = 0$) resulting in a period of free movement of belief. She decides to follow the signal and reject the investment (F). In spite of this, the following signal is incorrect and suggesting investing is optimal, so again she switches actions. Now, at point G, the second outer agent (player 2) also stops investing. When both outer players agree on the same action the period of free movement of belief has ended which forces player 1 to follow her neighbors ($\delta_{1t_{16}} = -2$). Irrespective of her signals, she adapts to the correct action.

This example illustrates the domino effect: after each change in the underlying state the group depends on the outer players that can freely alternate between actions to change and break the herd. The agent in the center is never the designated person to deviate. As soon as this agent observes a deviating action from a neighbor, she is free to change actions herself. In a line network, all nodes with two connections encounter this local gridlock effect during a herd. Until one of the outer players has decided to switch actions, the others act as a counterweight, opposing the altered distribution of signals resulting from the state change. Once an outer player has switched actions, her neighbor in the outer layer of inner nodes is free to choose the opposite action, and so forth. Proposition 3 describes the domino effect.

Proposition 3 *For an undirected line network, if actions have converged and* $1 \leq \alpha \leq 2\beta$*, agents with one link can adapt to the opposite action. They are expected to do so if* $p < \left(1 + \left(\frac{q}{1-q}\right)^{(1-2q)}\right)^{-1}$. *After these agents have changed, their neighbors can adapt as well. This process continues until all agents have adapted or the underlying state changes.*

These results extend to an undirected line network with any number of agents without loss of generality. Naturally, each added node means an additional agent with two neighbors. The domino effect describes that breaking the herd by switching actions will always start at the two players on the edge of the network and is expected to expand inward as long as the state remains fixed. The deviant actions taken by one of these agents reveals valuable information to others. The signals she received are opposing the herd and she is willing to follow these signals (Peck and Yang, 2011). In addition, we know that convergence of actions to an opposite state by all agents will take at least two periods for a three-person network. If the network is an n-person line network the transition to a herd on the opposite action will take at least $\frac{1}{2}n$ periods if n is even and $\frac{1}{2}(n+1)$ periods if n is uneven.

The domino effect can describe the adaptation process for any shape of undirected network that has a certain threshold comparable to (10). From the population of agents herding after the state of the world has changed, those with a limited number of connections are the ones that will have to fuel the adaptation process by switching first. Congestion(s) arise(s) where pockets of highly connected agents are located. As an example, consider a complex network composed of nodes with an unequal number of links in which all agents are taking the same action. Moreover, for all agents $\alpha = 2\frac{1}{2}\beta$ and (11) is satisfied. If these agents are herding, at first, adaptation to the underlying state is only possible for those with one or two neighbors. Agents with one neighbor are likely to change their action first, because they have only one observation counteracting their signals. Furthermore, if all agents that can adapt have done so, some agents might remain that still observe a difference of three neighbors taking the incorrect action ($\delta_{it} \geq 3$ or $\delta_{it} \leq -3$). These agents will never be able to adapt and the network will be divided between agents that adapt freely, agents that have to wait before a subset of their neighbors have changed and those that will never obtain a belief that will make them change actions (a local persistent herd). A fashion for some will be a fad for others.

6 Conclusion

The aim of the foregoing analysis is to understand the effect of various network architectures on the social learning process. Individuals receive a private signal and observe a subset of all actions taken in the previous period. The goal of each individual is to maximize utility by estimating what the true underlying time-varying state is. The results show that individual rational behavior may lead to collective inefficiencies, especially in highly connected societies.

Two effects of a dynamic state of the world have been identified. First, the *threat* of a change in underlying state ensures cascades are temporary. Information depreciates and the resulting revision of the belief parameter acts as an adaptation mechanism. Temporary cascades are a necessary condition for adaptation because they ensure agents have a continuous need for new information. Beliefs become active after some time and if the content of the information has changed this might lead them to try a different action. Furthermore, the presence of neighbors does not affect the evolution of belief during a cascade as individuals ignore all information. An increase in neighbors does decrease the likelihood that cascades occur, because individuals are less inclined to ignore a large group of neighbors taking the same action. Second, the *actual* change in the state of the world causes a change in the distribution of signals. Agents need to act when they notice a change in the signal distribution. The ability to adapt depends on the interplay between the reliability of neighbors, the quality of the private signal and the number of connections. An isolated agent, that relies solely on the private signal, will always be able to change its action in accordance with the state of the world. The expected signals adjust the belief of the agent towards the true state over time (similar to a sequential learning framework). Therefore, herds are likely as everyone is expected to choose the correct action, but herding is not a result of interaction and there is no fixed order of adaptation.

In a complete network I identify the gridlock effect as the cause of a persistent herd. Herding is permanent if everybody has multiple connections and the actions they observe are reliable. After actions have converged, the network takes over from the signal as the primary source of information, leading to a state of groupthink. People are ready to change to a different action but none of the agents is willing to deviate first. There is too much (incorrect) information to make a balanced decision, leading to a loss in utility for all agents.

Reducing the information available to the agents can alleviate the gridlock effect. In an incomplete network, time and place determine whether and when individuals are able to change actions. The domino effect describes that, at first, adaptation is possible for agents with a small number of connections. These agents will have to change first, while other highly connected individuals wait until they observe a different action from their neighbors. A change of action by sparsely connected agents is a message to others that they have received valuable new information about a possible change in the state of the world and they are willing to act accordingly. Change starting at the outskirts will then expand inward over time, provided that the state remains fixed during this process.

These findings are subject to several limitations.

7 Discussion

Here, I discuss three modelling assumptions and the way in which they affect the results.

The distribution of priors. The distribution of priors is the starting point for the evolution of belief. All agent have a flat prior, $\mu_{it_0} = \frac{1}{2}$, because both states of the world are equally likely at the start. As a consequence, there is a high probability that all agents take the correct action in the short-term. Furthermore, the presence of a network magnifies this effect as individuals that are incorrect initially will be corrected by their peers. As an example of a different distribution, Bala and Goyal (1998) assume each agent has a prior from a set of prior beliefs $\mu_{it_0}(\theta) \in \wp(\theta)$ (and to be complete these beliefs are interior $\mu_{it_0}(\theta) > 0$). Allowing for a diverse set of priors will change the outcome in the short-run. Convergence to the correct action is not self-evident, because some of the agents will choose incorrect actions without any information and can influence the others to join them in an incorrect herd. The long-term behavior and the adaptation process studied in the present paper are robust to this change. After actions converge, the flat prior is no longer important for the further adaptation process and the results follow. Furthermore, one of the advantages of a time-varying state is that an incorrect decision before a change is correct afterwards.

The subjective reliability of others. The Bayesian approach demands that the agent make inferences about the actions of others. For every observation, the subjective reliability of others (p) is this inference. As I have shown, the interpretation of this variable is troublesome and most of the time the agents are simply wrong about whether the actions of neighbors were correct. This imposed ignorance causes the persistent herd behavior found when neighbors are providing information. As an example, if agents herd on the incorrect action the actual value of the reliability of others is $p_t = P(z_{jt} = \theta_t) = 0$. None of the agents takes the correct decision in any of these periods and this is not common knowledge. As a consequence, information received from neighbors is overweighted.

A more realistic approach would be to let the subjective reliability of neighbors truly reflect the percentage of correct actions taken in the preceding period (without agents knowing which of their neighbors was correct). If the true reliability is announced publicly at the start of the period, this would gradually reduce the value of observations if these are incorrect, so that agents could infer that they have to switch actions to adapt. Hence, the reliability of neighbors is time-dependent, p_t. Moreover, as this process continues, and more people take the correct action, the confidence in neighbors will increase until beliefs are locked in their interval once again. Every change in the underlying state causes this process to recur.

The volatility of the dynamic state. It is assumed that the state of the world is not likely to change two periods in succession, $\varepsilon \in \left[0, \frac{1}{2}\right)$. This deserves two important (contradicting) remarks. First, the state should not be too volatile. The domino effect in an incomplete network can only be expected if the state is fixed for a number of periods. Beliefs in a cascade set need time to become active and once they are active it generally takes several periods before this leads to a change in action. Second, the state should change regularly. For $\varepsilon \approx 0$, the duration of temporary cascades can be longer than the duration of the game and there might not be enough opportunities to revise the belief for each agent. The mechanism that revises beliefs is then lacking and the results will only hold if time is infinite.

Chamley (2004) emphasizes the relevance of knowing at which speed learning occurs. The present paper is about (the absence of) an asymptotic result and not about the rate of learning. Nevertheless, the speed of learning is important as slow learning automatically implies slow adaptation. The adaptation mechanism needs to operate with sufficient speed to keep up with the ever changing state. The rate at which adaptation occurs is left for further research.

Appendix

Proof of Lemma 1: The starting point is (i) the threshold of the lower cascade set for an agent without neighbors (η_0). In the presence of neighbors (ii) a cascade starts when all observations and the signal are opposing the belief of the agent, but she is unwilling to alter the action. Finally, (iii) an increase in the number of neighbors decreases the cascade threshold.

(i) For an agent without neighbors and belief $\mu_{it} \leq \frac{1}{2}$ (taking action $x_{it} = 0$) the signal is ignored if the belief is in the lower cascade set, $\mu_{it} \in H_l = \left[0, \eta_0\right)$. Therefore, the threshold η_0 is the maximum belief for which inserting $s_{it} = 1$ in (4) will not lead to a change in action, $\mu_{it+1} \in \left[0, \frac{1}{2}\right]$. Remember that $\mu_{it+1} = \frac{1}{2} \Rightarrow \lambda_{it+1} = 0$, so that

$$\lambda_{it+1} = \ln\left(\frac{q}{1-q}\right) + \ln\left(\frac{\mu_{it}}{1-\mu_{it}}\right) = 0.$$

This gives

$$\mu_{it+1} = \eta_0 = 1 - q.$$

(ii) With observations, $|N_i| > 0$, the range of the cascade set depends on neighbors. For an agent with belief $\mu_{it} \leq \frac{1}{2}$ (taking action $x_{it} = 0$) a cascade has started if updating the belief will not change the action, $\lambda_{it+1} \leq 0$. She receives a signal contradicting her belief, $s_{it} = 1$, and observes all agents investing so that the observation sets are $Z^1_{it} = \{N_i\}$ and $Z^0_{it} = \{\emptyset\}$. Including this information in (4) means that $|\delta_{it}| = |N_i|$ and gives threshold $\eta_{|N_i|} = \left(1 + \frac{q}{1-q}\frac{p}{1-p}^{|N_i|}\right)^{-1}$. Similarly, $1 - \eta_{|N_i|} = \left(1 + \frac{1-q}{q}\frac{p}{1-p}^{-|N_i|}\right)^{-1}$ is the upper cascade threshold.

(iii) It is easy to see that $\eta_{|N_i|} < \eta_0$. Furthermore, $\eta_{|N_i|}$ decreases in $|N_i|$ for $p \in \left[\frac{1}{2}, 1\right]$ and $q \in \left(\frac{1}{2}, 1\right)$.■

Proof of Lemma 2: The first part of the proof is similar to that of Proposition 1 in Moscarini et al. (1998) and replicates these findings if

$N_i = \{\emptyset\}$. Here it is shown that (i) beliefs tend to the fixed point and (ii) there is a maximum duration of cascade belief.

(i) The evolution of belief during a cascade for $\varepsilon > 0$ ((7) in the text) can be written as

$$\mu_{it+1}(\theta_t) = \mu_{it}(1 - 2\varepsilon) + \varepsilon.$$

The belief is adjusted with constant value $|1 - 2\varepsilon|$ and $\dfrac{\partial \mu_{it+1}^d}{\partial \mu_{it}^d} = 1 - 2\varepsilon > 0$ for $\varepsilon \in \left(0, \frac{1}{2}\right)$. At the fixed point of this linear recurrence relation the prior is equal to the posterior. Therefore, $\mu_{it+1} = \mu_{it}$ implies $\mu_{it} = \mu_{it}(1 - 2\varepsilon) + \varepsilon$. It follows that $\mu_{it} = \frac{1}{2}$. As soon as a cascade starts, the belief of the agent will be directed towards this asymptote.

(ii) The lowest possible belief subject to a dynamic state is ε. Assume the agent has this belief and $\varepsilon < \eta$, so that the belief is in the cascade region. The belief increases each period following $\mu_{it} = \frac{1}{2} - \frac{1}{2}(1 - 2\varepsilon)^{t+1}$ if $\mu_{it} < \frac{1}{2}$ (and decreases as $\mu_{it} = \frac{1}{2} + \frac{1}{2}(1 - 2\varepsilon)^{t+1}$ if $\mu_{it} > \frac{1}{2}$). The agent returns to Bayesian updating as soon as $\frac{1}{2} - \frac{1}{2}(1 - 2\varepsilon)^{t+1} \geq \eta$, i.e. after a finite number of periods. For all other feasible cascade beliefs $\mu_{it} \in [\varepsilon, \eta)$ the cascade will end in less or an equal number of periods, as a consequence of the constant speed of movement towards the fixed point.■

Proof of Proposition 1: The proposition states that an isolated agent can adapt to a true state of the world. Lemma 2 describes that cascades are temporary if $\varepsilon > 0$ and the belief will exit the cascade region within finite time. Here, it is shown that (i) the distribution of signals ensures an isolated agent is expected to choose the correct action if the belief is outside the cascade interval, $\mu_{it} \in [\eta, 1 - \eta]$ (based on Chamley (2004)).

(i) Assume the agent invests ($x_{it} = 1$) but the true state is $\theta_t = 0$. The dynamic state ensures the belief is bounded: $\mu_{it} \in [\varepsilon, 1 - \varepsilon]$. Neighbors do not have influence, $E[\chi_{it}] = 0 \ \forall t$, and the expected value of the random variable is

$$E[\zeta_{it}|\theta_t = 0] = \ln\left(\frac{q}{1-q}\right)(1 - 2q).$$

Signals are informative $q \in \left(\frac{1}{2}, 1\right)$ so that $E\left[\zeta_{it}|\theta_t = 0\right] < 0$. Hence, the LLR tends to $-\infty$ and μ_{it} tends to ε as $t \to \infty$. If the state changes before $\mu_{it} \leq \frac{1}{2}$, the correct (and optimal) action is already chosen and $x_{it} = \theta_t$. If the state of the world is $\theta_t = 1$, the signal is again conditional on the true state of the world and

$$E\left[\zeta_{it}|\theta_t = 1\right] = \ln\left(\frac{q}{1-q}\right)(2q-1).$$

As $E\left[\zeta_{it}|\theta_t = 1\right] > 0$, the LLR tends to ∞ and μ_{it} tends to $1-\varepsilon$ as $t \to \infty$.■

Proof of Proposition 2: The proof consist of three parts. First, it is shown that (i) if (9) is satisfied none of the agents is able to switch actions once a herd has started. Furthermore, (ii) if there are two groups of equal size one of the agents will deviate and, as a result, (iii) actions will converge.

(i) During a persistent herd the signal is dominated by the information from observations. Signal precision has a constant value s.t. $|\zeta_{it}| = \ln\left(q/\left(1-q\right)\right)$. The influence of observations from neighbors depends on the number of neighbors agreeing on the same action, $|\chi_{it}| = \ln\left(p/\left(1-p\right)\right)^{|\delta_{it}|}$. By definition, in a herd all agents take the same action which implies $|\delta_{it}| = |N_i| = n-1$ for all i. Therefore, $|\chi_{it}| = \ln\left(p/\left(1-p\right)\right)^{|N_i|}$. The network dictates the direction of belief if

$$|\zeta_{it}| < |\chi_{it}| \Rightarrow \ln\left(\frac{q}{1-q}\right) < \ln\left(\frac{p}{1-p}\right)^{|N_i|}.$$

The threshold follows as

$$|N_i| > \frac{\alpha}{\beta}.$$

The LLR is (bounded and) limited to $(-\infty, 0)$ or $(0, \infty)$ and beliefs will be contained in their respective interval $\left[\varepsilon, \frac{1}{2}\right)$ or $\left(\frac{1}{2}, 1-\varepsilon\right]$ for $\varepsilon > 0$.

(ii) This is a proof by contradiction. Assume that society is divided into two groups of equal size, A and B, that continue to hold a different belief for the the rest of the game. Furthermore, during a finite period of time $T_k = (t_{k_0}, ..., t_{k_u})$, where $t_{k_0} \geq t_0$ and $t_{k_u} \leq T$, the state of the world is $\theta_{T_k} = 1$. At the start of t_{k_0}, all agents in group A chose to invest, with $\mu_{it_{k_0}} > \frac{1}{2}\ \forall i \in A$. The agents in group B are incorrect and believe that not investing is optimal, $\mu_{it_{k_0}} \leq \frac{1}{2}\ \forall i \in B$.

In any complete network $\delta_{it} \neq 0$ will always hold for one of the two groups. For groups that are split equally this implies $\delta_{it} = -1$ for all agents in group A and $\delta_{it} = 1$ for all agents in group B (none of the agents observes her own action as $i \notin N_i$). Any balance in opinion will shift as a result of the distribution of signals conditional on the state of the world. The LLR tends to ∞ if

$$E\left[\zeta_{it}|\theta_t\right] + \beta^{\delta_{it}} > 0.$$

Inserting the signal precision ratio gives

$$\alpha\left(2E\left[s_{it}|\theta_t\right] - 1\right) + \beta^{\delta_{it}} > 0.$$

Analogously,

$$\delta_{it} > \frac{\alpha}{\beta}\left(2\theta_t - 1\right)\left(1 - 2q\right) \qquad \text{for } \theta_t \in \{0, 1\}.$$

If this does not hold $E\left[\lambda_{it+1}\right] < 0$, and the LLR tends to $-\infty$. A situation in which group A and B continue to engage in different actions can only be sustained as long as the LLR of group A ($\delta_{it} = -1$) tends to ∞ and the LLR of group B ($\delta_{it} = 1$) tends to $-\infty$:

$$-1 > \frac{\alpha}{\beta}\left(2\theta_t - 1\right)\left(1 - 2q\right) > 1.$$

This is contradiction as $(1 - 2q) \in (-1, 0)$ and $\alpha/\beta > 0$ for $\frac{1}{2} > p > 1$ and $\frac{1}{2} > q > 1$.

(iii) Actions converge because, provided that the state is fixed during a certain period, it is not possible that one or more individuals continue to hold a different belief than the rest of the group as shown by (ii). At some moment in time an agent in group B is expected to change. These individuals will be drawn towards the correct side, because the signal and the observations indicate their decision is incorrect and the belief will be adjusted upward. For the agents in group A their signal indicates they are correct, but their observation indicates otherwise. Whether their belief is directed downward depends on the strength of the signal. Eventually, one of the agents switches from A to B, or vice versa. This switch will accelerate the convergence process.

Assume one of the agents changes her decision from $x_{it} = 0$ to $x_{it+1} = 1$. The content of the observation set for this agent will not change (she does not observe her own action), but her belief is already $\mu_{it+1} > \frac{1}{2}$ and $\delta_{it+1} = 1 \Rightarrow E[\lambda_{it+1}] > 0$. For the remaining agents in group A their prejudice has been confirmed and $\delta_{it+1} = 1 \Rightarrow E[\lambda_{it+1}] > 0$. Their belief now tends to $1 - \varepsilon$. Agents in group B observe one less companion and an additional dissident, so that $\delta_{it+1} = 3$. They will eventually follow the agent that has first switched to group B. As soon as $|\delta_{it}| > \alpha/\beta$ for A or B, the network takes over as primary source of information.■

Proof of Proposition 3: It is assumed that actions have converged. The domino effect will only occur if (i) the threshold (10) is satisfied. Furthermore, the signal is expected to correct beliefs (11) if (ii) agents collectively invests and the true state is $\theta_t = 0$ and (iii) agents collectively abstain and the true state is $\theta_t = 1$.

(i) The threshold for the domino effect is found by inserting $|N_i| = 1$ for the outer nodes and $|N_i| = 2$ for the inner nodes in (9).

(ii) The expected distribution of signals directs the belief downward if $E[\lambda_{it}] < 0$. If the outer node has belief $\mu_{it} > \frac{1}{2}$ and observes $\delta_{it} = 1$, the LLR tends to $-\infty$ if

$$E[\zeta_{it}|\theta_t = 0] + \chi_{it} < 0,$$

so that

$$\ln\left(\frac{q}{1-q}\right)(1-2q)+\ln\left(\frac{p}{1-p}\right)<0.$$

This gives

$$\left(\frac{q}{1-q}\right)^{(1-2q)}<\frac{1-p}{p}\Rightarrow p<\left(1+\left(\frac{q}{1-q}\right)^{(1-2q)}\right)^{-1}. \tag{A1}$$

(iii) If no one invests and the state is $\theta_t = 1$, the expected distribution of signals directs the belief upward if $E[\lambda_{it}] > 0$. If the outer node has belief $\mu_{it} \leq \frac{1}{2}$ and observes $\delta_{it} = -1$, the LLR tends to ∞ if

$$\ln\left(\frac{q}{1-q}\right)(2q-1)+\ln\left(\frac{p}{1-p}\right)^{-1}>0.$$

Accordingly,

$$\left(\frac{q}{1-q}\right)^{(2q-1)}>\frac{p}{1-p}\Rightarrow p<\left(1+\left(\left(\frac{q}{1-q}\right)^{(2q-1)}\right)^{-1}\right)^{-1}. \tag{A2}$$

Here (A1) and (A2) are equivalent. If these equations hold the single counteracting observation is overruled by the expected sequence of signals.■

References

Acemoglu, D., Dahleh, M.A., Lobel, I., Ozdaglar, A. (2011). Bayesian Learning in Social Networks, *Review of Economic Studies*, 78 (4), 1201-1236.

Bala, V., Goyal, S. (1998). Learning from Neighbours, *Review of Economic Studies*, 65 (3), 595-621.

Banerjee, A.V. (1992). A Simple Model of Herd Behavior, *The Quarterly Journal of Economics*, 107 (3), 797-817.

Bikhchandani, S., Hirshleifer, D., Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades, *Journal of Political Economy*, 100 (5), 992-1026.

Bikhchandani, S., Hirshleifer, D., Welch, I. (2008). Information Cascades. In Durlauf, S.N., Blume, L.E. (Eds.) *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, doi:10.1057/9780230226203.0799.

Çelen, B., Kariv, S. (2004). Distinguishing Informational Cascades from Herd Behavior in the Laboratory, *The American Economic Review*, 94 (3), 484-498.

Chamley, C.P. (2004). *Rational Herds: Economic Models of Social Learning*. Cambridge: Cambridge University Press.

Chandrasekhar, A.G., Larreguy, H., Xandri, J.P. (2015). Testing Models of Social Learning on Networks: Evidence from a Lab Experiment in the Field, NBER Working Paper, No. 21468.

DeGroot, M.H. (1974). Reaching a Consensus, *Journal of the American Statistical Association*, 69 (345), 118-121.

DeMarzo, P.M., Vayanos, D., Zwiebel, J. (2003). Persuasion Bias, Social Influence, and Unidimensional Opinions, *The Quarterly Journal of Economics*, 118 (3), 909-968.

Ellison, G., Fudenberg, D. (1993). Rules of Thumb for Social Learning, *Journal of Political Economy*, 101 (4), 612-643.

Ellison, G., Fudenberg, D. (1995). Word-of-mouth Communication and Social Learning, *The Quarterly Journal of Economics*, 110 (1), 93-125.

Enke, B., Zimmermann, F. (2013). Correlation Neglect in Belief Formation, IZA Discussion Paper, No. 7372.

Eyster, E., Weizsäcker, G. (2011). Correlation Neglect in Financial Decision-Making, DIW Berlin Discussion Paper, No. 1104.

Frongillo, R.M., Schoenebeck, G., Tamuz, O. (2011). Social Learning in a Changing World, The 7th Workshop on Internet and Network Economics 2011, arXiv:1109.5482 [cs.SI].

Gale, D., Kariv, S. (2003). Bayesian Learning in Social Networks, *Games and Economic Behavior*, 45 (2), 329-346.

Goeree, J.K., Palfrey, T.R., Rogers, B.W. (2006). Social Learning with Private and Common Values, *Economic Theory*, 28 (2), 245-264.

Goeree, J.K., Palfrey, T.R., Rogers, B.W., McKelvey, R.D. (2007). Self-Correcting Information Cascades, *Review of Economic Studies*, 74 (3), 733-762.

Golub, B., Jackson, M.O. (2010). Naïve Learning in Social Networks and the Wisdom of Crowds, *American Economic Journal: Microeconomics*, 2 (1), 112-149.

González-Avella, J.C., Eguíluz, V.M., Marsili, M., Vega-Redondo, F., San Miguel, M. (2011). Threshold Learning Dynamics in Social Networks, *PLoS One*, 6 (5), e20207.

Goyal, S. (2007). *Connections: An Introduction to the Economics of Networks.* Princeton: Princeton University Press.

Jackson, M.O. (2008). *Social and Economic Networks.* Princeton: Princeton University Press.

Kanoria, Y., Tamuz, O. (2013). Tractable Bayesian Social Learning on Trees, *IEEE Journal on Selected Areas in Communications*, 31 (4), 756-765.

Klumpp, T. (2006). Linear Learning in Changing Environments, *Journal of Economic Dynamics and Control*, 30 (12), 2577-2611.

Lamberson, P.J. (2010). Social Learning in Social Networks, *The B.E. Journal of Theoretical Economics*, 10 (1), Article 36.

Levy, G., Razin, R. (2015). Correlation Neglect, Voting Behavior and Information Aggregation, *The American Economic Review*, 105 (4), 1634-1645.

Moscarini, G., Ottaviani, M., Smith, L. (1998). Social Learning in a Changing World, *Economic Theory*, 11 (3), 657-665.

Orléan, A. (1995). Bayesian Interactions and Collective Dynamics of Opinion: Herd Behavior and Mimetic Contagion, *Journal of Economic Behavior and Organization*, 28 (2), 257-274.

Peck, J., Yang, H. (2011). Investment Cycles, Strategic Delay, and Self-Reversing Cascades, *International Economic Review*, 52 (1), 259-280.

Shahrampour, S., Rakhlin, A., Jadbabaie, A. (2013). Online Learning of Dynamic Parameters in Social Networks, Advances in Neural Information Processing Systems 2013, arXiv:1310.0432 [math.OC].

Smith, L., Sørensen, P. (2000). Pathological Outcomes of Observational Learning, *Econometrica*, 68 (2), 371-398.

Tadelis, S. (2013). *Game Theory: An Introduction.* Princeton: Princeton University Press.