

Contagious Synchronization and Endogenous Network Formation in Financial Networks *

Christoph Aymanns[†], Co-Pierre Georg[‡]

July 1, 2014

Abstract

When banks choose similar investment strategies the financial system becomes vulnerable to common shocks. We model a simple financial system in which banks decide about their investment strategy based on a private belief about the state of the world and a social belief formed from observing the actions of peers. Observing a larger group of peers conveys more information and thus leads to a stronger social belief. Extending the standard model of Bayesian updating in social networks, we show that the probability that banks synchronize their investment strategy on a state non-matching action critically depends on the weighting between private and social belief. This effect is alleviated when banks choose their peers endogenously in a network formation process, internalizing the externalities arising from social learning.

*We would like to thank Toni Ahnert, Jean-Edouard Colliard, Jens Krause, Tarik Roukny, three anonymous referees, seminar participants at ECB, Bundesbank, as well as the 2013 INET Plenary Conference in Hong Kong, and the VIII Financial Stability Seminar organized by the Banco Central do Brazil for helpful discussions and comments. This paper has been prepared under the Lamfalussy Fellowship Program sponsored by the ECB whose support is gratefully acknowledged. The views expressed in this paper do not necessarily reflect the views of Deutsche Bundesbank, the ECB, or the ESCB.

[†]Mathematical Institute, University of Oxford and Institute for New Economic Thinking at the Oxford Martin School. E-Mail: aymanns@maths.ox.ac.uk.

[‡]Deutsche Bundesbank and University of Cape Town Graduate School of Business. E-Mail: co-pierre.georg@bundesbank.de.

Keywords: social learning, endogenous financial networks, multi-agent simulations, systemic risk

JEL Classification: G21, C73, D53, D85

1 Introduction

When a large number of financial intermediaries choose the same investment strategy (i.e. their portfolios are very similar) the financial system as a whole becomes vulnerable to common shocks. A case at hand is the financial crisis of 2007/2008 when many banks invested into mortgage backed securities in anticipation that the underlying mortgages—many of which being US subprime mortgages—would not simultaneously depreciate in value. This assumption turned out to be incorrect resulting in one of the largest financial crises since the great depression. How could so many banks choose a non-optimal investment strategy despite the fact that they carefully monitor both economic fundamentals and the actions of other banks?

This paper presents a simple agent-based model in which financial intermediaries synchronize their investment strategy on a state non-matching action despite informative private signals about the state of the world. In a countable number of time-steps N agents, representing financial intermediaries (banks for short), choose one of two actions. There are two states of the world which are revealed at the end of the simulation. A bank's action is either state-matching, in which case the bank receives a positive payoff if the state is revealed, or it is state-non-matching in which case the bank receives zero. Banks are connected to a set of peers in a financial network of mutual lines of credit resembling the interbank market. They receive a private signal about the state of the world and observe the previous actions of banks with whom they are connected via a mutual line of credit, but not of other banks. Based on both signals banks form a belief about the state of the world and choose their action accordingly.

Our model differs from the existing literature along two dimensions. First and foremost we develop an agent-based model of the financial system with strategic interaction amongst agents. This differs from existing models (see,

for example, Poledna et al. (2014), Bluhm et al. (2013), Georg (2013), and Ladley (2013)) where agent behaviour is myopic. Agents in myopic models react to the state of the world but when choosing an optimal action they do not take into account how other agents will react to their choice. Thus, the notion of equilibrium in myopic models is a rather mechanical. Strategic interaction amongst agents arises in our model from the fact that agents learn about their neighbors' actions, i.e. via the social belief. All the aforementioned papers furthermore use an exogenous network structure as starting point for the agent-based simulation, while our model uses an endogenous network formation process to arrive at a pairwise stable network structure that maximizes expected utility from social learning.

Second, while our model is mildly boundedly rational it shares a number of assumptions with the literature on Bayesian learning in social networks. The main difference to this literature (see, for example, Acemoglu et al. (2011), Gale and Kariv (2003)) is that we model an externality that is not present in the standard model of Bayesian learning in social networks. We assume that banks receive more information about the actions of other banks than they can computationally use. This assumption seems natural in a financial system that is increasingly complex.¹ The underlying assumption is that banks cannot adjust their actions (i.e. their investment strategy) as fast as they receive information from their peers and thus have to aggregate over potentially large amounts of information.² The social belief in our model is formed not just from observing one neighbor at a time, but rather from observing a set of neighbors simultaneously. It is thus reasonable to assume that receipt of more information, i.e. observing the actions of a larger subset of agents, will create a stronger social belief than the receipt of less information. We model this by allowing for different weights of the social

¹See, for example, Haldane and Madouros (2012) for a discussion of increasing complexity in financial regulation.

²This assumption renders the agents boundedly rational, albeit mildly so, as for example DeMarzo et al. (2003) argue.

and private belief.

The other key difference to the existing literature on Bayesian learning in social networks is that we allow agents to endogenously form links based on the utility they get from an improved social belief in a first stage of the model. In the second stage of the model agents then learn about the state of the world and take their investment decisions. To the best of our knowledge, the only other paper considering endogenously formed social networks in a Bayesian learning setup is Acemoglu et al. (2014) who develop a two-stage game similar to ours. Acemoglu et al. (2014) model endogenous network formation via a communication cost matrix where some agents (in a *social clique*) can communicate at low costs, while others communicate at high cost. The main difference to our model is that we endogenously obtain a decreasing marginal value of additional links. We obtain a resulting endogenous network structure which is pairwise stable in the sense of Jackson and Wolinsky (1996).

We obtain two sets of results, one for networks with exogenous network structure, and one for endogenously formed networks. First, we analyze different ways of weighting private and social belief. In particular, we compare the standard *equal weighting* scenario in which agents place equal weights on their private and social belief, with two scenarios where agents place more weight on the social belief when they have a larger neighborhood. In the *neighborhood size* scenario the social belief is weighted with the size of the neighborhood, i.e. the private signal is weighted equal to every observed neighbor action. In the *relative neighborhood* scenario agents put more weight on the social belief when the neighborhood constitutes a larger share of the overall network. For completely uninformative signals there is no difference between these weighting functions. For informative signals, however, the weighting function has an impact on the probability that agents synchronize their investment decisions on a state non-matching ac-

tion, i.e. for the probability that choosing a state-non-matching action is contagious.³ Contagion is, very generally, understood as the transmission of adverse effects from one agent to another and is more likely if agents place greater weight on their social belief and depends on the density of the underlying exogenous network structure.⁴ The probability of contagion increases by a factor of 200 in the neighborhood size scenario compared to the equal weighting scenario, which highlights the importance of understanding the learning dynamics when agents place different weight on their social belief, depending on the size of their neighborhood.

We show that contagious synchronization occurs even if private signals are informative and if agents are initialized with an action that is on average state matching. The probability of contagion depends non-monotonously on the density of the network. For small network densities $\rho \lesssim 0.1$ the probability of contagion increases sharply and then decreases slowly for larger network densities. We confirm the robustness of our results by conducting 2,000 independent simulations where we observe the average final action as a function of the average initial action with varying network densities.

This result is of particular interest for policy makers as it relates two sources of systemic risk: common shocks and interbank market freezes. When the network density is too small, for example in the aftermath of an interbank market freeze, banks are unable to fully incorporate the information about their peers' actions. This effect is empirically tested by Caballero (2012), who documents a higher correlation amongst various asset classes in the world in the aftermath of the Lehman insolvency, i.e. during times of extreme stress on interbank markets and heightened uncertainty about the state of the world. This can be understood as a contagious synchronization

³Such *informational cascades* are a well-documented empirical phenomenon. See, for example, Alevy et al. (2007), Bernhardt et al. (2006), Chang et al. (2000), Chiang and Zheng (2010), and Cipriani and Guarino (2014).

⁴For a more thorough discussion of the different forms of contagion, see for example Bandt et al. (2009).

of bank's investment strategies for which our model provides a simple rationale.

Second, turning to the extension of endogenously formed networks, we show that endogenous link formation in the first stage of our model can significantly improve the speed of learning and reduce the probability of contagious synchronization relative to random networks. When private signals are less informative, the additional utility from forming a link is smaller and the endogenously formed network is less dense. This in turn can increase the probability of contagious synchronization in the second stage of the model. Heightened uncertainty about the state of the world, i.e. a less informative signal, does therefore not only directly increase the probability of contagious synchronization, but also indirectly because agents have less incentives to endogenously form links. If agents are heterogeneous in the informativeness of their private signals, i.e. if some agents receive signals with higher precision than others, we show that the resulting endogenous network structure is of a core-periphery type. The structure of real-world interbank markets is often of this particular type, as for example Craig and von Peter (2014) show. Naturally, these endogenously formed networks transfer information more effectively from highly informed agents to less informed agents than simple random networks.

This paper relates to three strands of literatures. First and foremost, the paper develops a financial multi-agent simulation in which agents learn not only from private signals, but also via endogenously formed interbank links. This is in contrast with existing multi-agent models of the financial system which include Nier et al. (2007) and Iori et al. (2006) who take a fixed network and static balance sheet structure.⁵ Slight deviations from these

⁵Closely related is the literature on financial networks. See, for example, Allen and Gale (2000), and Freixas et al. (2000) for an early model of financial networks. The vast majority of models in this literature consider a fixed network structure only (see, amongst various others, Battiston et al. (2012)).

models can be found, for example, in Bluhm et al. (2013), Ladley (2013), and Georg (2013) who employ different equilibrium concepts. The main contribution this paper makes is to develop a sufficiently simple model of a financial system with a clear notion of equilibrium that allows to be implemented on a computer and tested against analytically tractable special cases.

Second, this paper relates to the literature on endogenous network formation pioneered by Jackson and Wolinsky (1996) and Bala and Goyal (2000). Few papers on endogenous network formation in interbank markets, exist, however. Notable exceptions are Castiglionesi and Navarro (2007) who study the formation of endogenous networks in a banking network with microfounded banking behaviour. Unlike Castiglionesi and Navarro (2007), however, our paper uses a starkly simplified model of social learning to describe the behaviour of banks. This allows the introduction of informational spillovers from one bank to another, a mechanism not present in the work of Castiglionesi and Navarro (2007).

Finally, our paper is closely related to the literature on Bayesian learning in social networks. The paper closest to ours in this literature is Acemoglu et al. (2014) who study a model of sequential learning in an endogenously formed social network where each agent receives a private signal about the state of the world and observe past actions of their neighbors. We contribute to this literature by allowing agents to place more weight on their social belief when their neighborhood is larger. Other related papers in this literature include Banerjee (1992), Bikhchandani et al. (1992), Bala and Goyal (1998), and Gale and Kariv (2003) who, however, all consider static networks only.

The remainder of this paper is organized as follows. The next section develops the baseline model and presents the results in the limiting case of an

exogenous network structure. Section 3 generalizes the model by allowing agents to endogenously form links in a first stage of the model. Section (4) concludes.

2 Contagious Synchronization with Fixed Network Structure

2.1 Model Description and Timeline

There is a countable number of dates $t = 0, 1, \dots, T$ and a fixed number $i = 1, \dots, N$ of agents A^i which represent financial institutions and are called banks for short. By a slight abuse of notation the model parameter θ is sometimes called the state of the world and we assume it can take two values $\theta \in \{0, 1\}$. The probability that the world is in state θ is denoted as $\mathbb{P}(\theta)$ and we assume that each state of the world is obtained with equal probability $\frac{1}{2}$. At each point in time t bank i chooses one of two investment strategies $x_t^i \in \{0, 1\}$ which yields a positive return if the state of the world is revealed and matches the investment strategy chosen, and nothing otherwise. Agents take an action by choosing a certain investment strategy. Taking an action and switching between actions is costless. The utility of bank i from investing is given as:

$$u^i(x^i, \theta) = \begin{cases} 1 & \text{if } x^i = \theta \\ 0 & \text{else} \end{cases} \quad (1)$$

The state of the world is unknown ex-ante and revealed at time T . This setup captures a situation where the state of the world is revealed less often (e.g. quarterly) than banks take investment decisions (e.g. daily).⁶

⁶In an alternative setup the state of the world is fixed throughout and an agent collects information and takes an irreversible decision at time t , but receives a payoff that is discounted by a factor $e^{-\kappa t}$. Both formulations incentivize agents to take a decision in finite time instead of collecting information until all uncertainty is eliminated.

Banks can form interconnections in the form of mutual lines of credit. The set of banks is denoted $N = \{1, 2, \dots, n\}$ and the set of banks to which bank i is directly connected is denoted $K^i \subseteq N$. Bank i thus has $k^i = |K^i|$ direct connections called neighbors. This implements the notion of a network of banks \mathbf{g} which is defined as the set of banks together with a set of unordered pairs of banks called (undirected) links $L = \cup_{i=1}^n \{(i, j) : j \in K^i\}$. A link is undirected since lines of credit are mutual and captured in the symmetric adjacency matrix g of the network. Whenever a bank i and j have a link, the corresponding entry $g^{ij} = 1$, otherwise $g^{ij} = 0$. When there is no risk of confusion in notation, the network \mathbf{g} is identified by its adjacency matrix g . For the remainder of this section, I assume that the network structure is exogenously fixed and does not change over time. I assume that banks monitor each other continuously when granting a credit line and thus observe their respective actions.

In this section, the network \mathbf{g} is exogenously fixed throughout the simulation. In $t = 0$ there is no previous decision of agents. Thus, each bank decides on its action in autarky. Banks receive a signal about the state of the world and form a private belief upon which they decide about their investment strategy $x_{t=0}^i$. The private signal received at time t is denoted $s_t^i \in \overline{S}$ where \overline{S} is a Euclidean space. Signals are independently generated according to a probability measure \mathbb{F}_θ that depends on the state of the world θ . The signal structure of the model is thus given by $(\mathbb{F}_0, \mathbb{F}_1)$. I assume that \mathbb{F}_0 and \mathbb{F}_1 are not identical and absolutely continuous with respect to each other. Throughout this paper I will assume that \mathbb{F}_0 and \mathbb{F}_1 represent Gaussian distributions with mean and standard deviation (μ_0, σ_0) and (μ_1, σ_1) respectively.

In $t = 1, \dots$ bank i again receives a signal s_t^i but now also observes the $t - 1$ actions x_{t-1}^j of its neighbors $j \in K^i$. The model outlined in this section is implemented in a multi-agent simulation where banks are the agents. Date

$t = 0$ in the model timeline is the initialization period. Subsequent dates $t = 1, \dots, T$ are the update steps which are repeated until the state of the world is being revealed in state T . Once the state is revealed, returns are realized and measured. In the simulation results discussed in Section 2.2 the state of the world was revealed at the end of the simulation after the system has reached a steady state in which agents do not change their actions any more.⁷

Banks form a private belief at time t based on their privately observed signal s_t^i and a social belief based on the observed actions x_{t-1}^j their neighboring banks took in the previous period. The first time banks choose an action is a special case of the update step with no previous decisions being taken. The information set I_t^i of a bank i at time t is given by the private signal s_t^i , the set of banks connected to bank i in $t - 1$, K_{t-1}^i , and the actions x_{t-1}^j of connected banks $j \in K_{t-1}^i$. Formally:

$$I_t^i = \{s_t^i, K_{t-1}^i, x_{t-1}^j \forall j \in K_{t-1}^i\} \quad (2)$$

The set of all possible information sets of bank i is denoted by \mathcal{I}^i . A strategy for bank i selects an action for each possible information set. Formally, a strategy for bank i is a mapping $\sigma^i : \mathcal{I}^i \rightarrow x^i = \{0, 1\}$. The notation $\sigma^{-i} = \{\sigma^1, \dots, \sigma^{i-1}, \sigma^{i+1}, \sigma^n\}$ is used to denote the strategies of all banks other than i .

A strategy profile $\sigma = \{\sigma^i\}_{i \in 1, \dots, n}$ is a pure strategy equilibrium of this game of social learning for a bank i 's investment, if σ^i maximizes the bank's expected pay-off, given the strategies of all other banks σ^{-i} . Acemoglu et al.

⁷In practice this is ensured by having many more update steps than it takes the system to reach a steady state.

(2011) show that the strategy decision of bank i , $x_t^i = \sigma^i(I_t^i)$ is given as:

$$x^i = \begin{cases} 1 & \text{if } \mathbb{P}_\sigma(\theta = 1|s_t^i) + \mathbb{P}_\sigma(\theta = 1|x_{t-1}^j, j \in K_{t-1}^i) > 1 \\ 0 & \text{if } \underbrace{\mathbb{P}_\sigma(\theta = 1|s_t^i)}_{\text{private belief } p} + \underbrace{\mathbb{P}_\sigma(\theta = 1|x_{t-1}^j, j \in K_{t-1}^i)}_{\text{social belief } q} < 1 \end{cases} \quad (3)$$

and $x^i \in \{0, 1\}$ otherwise. The first term on the right-hand side of Equation 3 is the private belief, the second term is the social belief, and the threshold is fixed to $\frac{1}{2}$. This equation can be generalized when introducing weights on the private and social belief. In its most general form, it can be written as:

$$x^i = \begin{cases} 1 & \text{if } t(p^i, q^i) > \frac{1}{2} \\ 0 & \text{if } t(p^i, q^i) < \frac{1}{2} \end{cases} \quad (4)$$

where $t(p^i, q^i)$ is a weighting function depending on the private and social belief. A simple weighting function

$$t(p^i, q^i) = \begin{cases} \frac{1}{2}(p^i + q^i) & \text{if } |K^i| > 0, \\ p^i & |K^i| = 0, \end{cases} \quad (5)$$

implements the model of Acemoglu et al. (2011) where agents place equal weight on their private and social belief (we denote this weighting function as the *equal weighting* scenario).

The simple weighting function is appropriate in a setting where agents receive two signals at a time only: their private signal and the signal of their direct predecessor in the social network. In our setting, however, banks receive multiple signals, only one of which is their private signal. It is thus natural to allow for more general weighting functions which, however, have to satisfy two conditions: (i) In the case with no social learning ($K_{t-1}^i = \emptyset$), the weighting should reduce to the simple case $t(p^i, q^i) = p^i$ in which the agent will select action $x^i = 1$ whenever it is more likely that the state of

the world is $\theta = 1$ and zero otherwise; and (ii) With social learning the weighting should depend on the number of neighboring signals, i.e. the size of the neighborhood $k_{t-1}^i = |K_{t-1}^i|$. The underlying assumption is that the agent will place a higher weight on the social belief when the neighborhood is larger. We consider two different scenarios for the weighting function: (i) The private signal and each observed action are equally weighted (called the *neighborhood size* scenario):

$$t(p^i, q^i) = \left(\frac{1}{k_{t-1}^i + 1} \right) p^i + \left(\frac{k_{t-1}^i}{k_{t-1}^i + 1} \right) q^i \quad (6)$$

And (ii) Observed actions are weighted with the relative size of the neighborhood (called the *relative neighborhood* scenario):

$$t(p^i, q^i) = \left(1 - \frac{k_{t-1}^i}{N-1} \right) p^i + \left(\frac{k_{t-1}^i}{N-1} \right) q^i \quad (7)$$

The private belief of bank i is denoted $p^i = \mathbb{P}(\theta = 1 | s^i)$ and can easily be obtained using Bayes' rule. It is given as:

$$p^i = \left(1 + \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s_t^i) \right)^{-1} = \left(1 + \frac{f_0(s_t^i)}{f_1(s_t^i)} \right)^{-1} \quad (8)$$

where f_0 and f_1 are the densities of \mathbb{F}_0 and \mathbb{F}_1 respectively. Bank i is assumed to form a social belief q^i by simply averaging over the actions of all neighbors $j \in K_{t-1}^i$:

$$q^i = \mathbb{P}_\sigma(\theta = 1 | K_t^i, x^j, j \in K_{t-1}^i) = 1/k_{t-1}^i \sum_{j \in K_{t-1}^i} x_{t-1}^j \quad (9)$$

Given these private and social beliefs, agents choose an action according to equation (4).

Averaging over the actions of neighbors is a special case of DeGroot (1974) who introduces a model where a population of N agents is endowed with initial opinions $p(0)$. Agents are connected to each other but with varying

levels of trust, i.e. their interconnectedness is captured in a weighted directed $n \times n$ matrix T . A vector of beliefs p is updated such that $p(t) = Tp(t-1) = T^t p(0)$. DeMarzo et al. (2003) point out that this process is a boundedly rational approximation of a much more complicated inference problem where agents keep track of each bit of information to avoid a persuasion bias (effectively double-counting the same piece of information). Therefore, the model this paper develops is also boundedly rational.⁸

The model in this section can be formulated as an agent-based model. Banks are agents a^i who choose one of two actions $x^i \in \{0, 1\}$. Variables determined in the model internally are given by the private and social belief of agent i at time t , p_t^i, q_t^i and the only exogenously given parameter is the state of the world θ which is identical for all agents i . Each agent has an information set I^i given by Equation (2). Agents receive utility (1) and decide on their optimal strategy given in Equation (3). The interaction of agents is captured in a network structure \mathbf{g} , encapsulated in an agent i 's information set. Equations (8) and (9) specify how agents take their decisions and choose an optimal strategy.

2.2 Herding with Exogenous Network Structures

Our interest is to understand under which conditions agents in the model with an exogenously fixed network structure coordinate on a state non-matching action. Before analyzing the full model, we build some intuition by discussing useful benchmark cases. Let the state of the world be $\theta = 0$ and assume that $f_0 = mf_1$. For $m = 1$ the signal is completely uninformative. For $m > 1$ the signal is informative and more so the larger m is. In the equal weighting scenario, equation (3) together with equations (8) and (9)

⁸This bounded rationality can be motivated analogously to DeMarzo et al. (2003) who argue that the amount of information agents have to keep track of increases exponentially with the number of agents and increasing time, making it computationally impossible to process all available information.

yields:

$$\frac{1}{2}(1+m)^{-1} + \frac{1}{2k} \sum_{j \in K_{t-1}^i} x_{t-1}^j > \frac{1}{2} \Leftrightarrow \sum_{j \in K_{t-1}^i} x_{t-1}^j > k \left[1 - \frac{1}{(1+m)} \right] \quad (10)$$

For completely uninformative signals, $m = 1$, equation (10) implies that an agent will always follow the majority of her neighbors. For very informative signals, $m \gg 1$, equation (10) implies that an agent will only ignore her private signal if she receives a strong social signal. The required strength of the social signal increases with the precision of the private signal. Now consider the neighborhood size scenario. The equation analogous to (10) for this scenario reads:

$$\left(\frac{1}{k_{t-1}^i + 1} \right) (1+m)^{-1} + \left(\frac{k_{t-1}^i}{k_{t-1}^i + 1} \right) \frac{1}{k_{t-1}^i} \sum_{j \in K_{t-1}^i} x_{t-1}^j > \frac{1}{2} \quad (11)$$

$$\Leftrightarrow \sum_{j \in K_{t-1}^i} x_{t-1}^j > \frac{k_{t-1}^i}{2} + \left[\frac{1}{2} - \frac{1}{1+m} \right] \quad (12)$$

For completely uninformative signals $m = 1$ this condition reduces to the equal weighting scenario. For highly informative signals, $m \gg 1$, however, the agent is almost as willing to follow her neighbors as in the uninformative equal weighting scenario. The neighborhood scenario thus captures the situation where the agent is aware not only of her own private signal informativeness, but also of that of her neighbors. In the relative neighborhood scenario equation (10) reads:

$$\left(1 - \frac{k_{t-1}^i}{N-1} \right) (1+m)^{-1} + \left(\frac{k_{t-1}^i}{N-1} \right) \frac{1}{k_{t-1}^i} \sum_{j \in K_{t-1}^i} x_{t-1}^j > \frac{1}{2} \quad (13)$$

$$\Leftrightarrow \sum_{j \in K_{t-1}^i} x_{t-1}^j > (N-1) \left[\frac{1}{2} - \frac{1}{1+m} \right] + \frac{k_{t-1}^i}{1+m} \quad (14)$$

Again, for completely uninformative signals this equation reduces to the equal weighting case. For highly informative signals, $m \gg 1$, this equation reduces to $\sum_{j \in K_{t-1}^i} x_{t-1}^j > \frac{1}{2}(N-1)$ if the neighborhood is sufficiently

small $m \gg k_{t-1}^i$. An agent will thus ignore her private signal and follow the majority of her neighbors, if her neighborhood is larger than half of the network. The central node in a star network will thus always follow the majority of the spokes and spokes will always follow their private signal.

We can gain further insights into the model dynamics by resorting to a mean-field approximation in which we consider the simplified action dynamics of a representative agent. Given the adjacency matrix g of a network \mathbf{g} , the social belief q of the representative agent is given as $q = g\mathbf{x}/k$ where \mathbf{x} is the vector of all agent's actions. The social belief in the mean-field approximation is simply the average action of the population:

$$q = \Pr(x = 0 \mid q = q) \cdot 0 + \Pr(x = 1 \mid q = q) \cdot 1 = \Pr(x = 1 \mid q = q) \quad (15)$$

which yields a self-consistency relation for the social belief and hence for the average action of the population. The equilibrium average action q^* is implicitly given by the solution to the self-consistency condition:

$$q^* = \Pr(x = 1 \mid q = q^*) = \int_{1-q^*}^1 p_P(p \mid \theta = 0) dp. \quad (16)$$

Note that $\Pr(x = 1 \mid q = 0) = 0$ and $\Pr(x = 1 \mid q = 1) = 1$. Since $\Pr(x = 1 \mid q = q^*)$ is the cumulative distribution function of the bell shaped private belief it will have a sigmoid shape. Therefore the self consistency equation will have three solutions: $q_1 = 0$, $q_2 = 1$ and some $q_3 \in (0, 1)$. We illustrate this in figure 4. q_1 , q_2 are stable fixed points while q_3 is unstable (this can be seen graphically in figure 4 and is a direct result from the specified learning dynamics). q_3 defines the “critical” social belief beyond which the system synchronizes on the state non-matching action.

These exercises show that the impact of the weighting function on agents' strategies is easily understood in the case of completely uninformative and fully informative signals or in a mean field approximation. But what hap-

pens in the more realistic interim region? Are densely connected networks more conducive for agents to coordinate on state non-matching actions or sparse networks? And how does the fraction of nodes that coordinate on a state non-matching action depend on the initial conditions? We address these questions in an agent-based simulation for three cases. In the case (I) of *informed agents* the distance between the mean of the two signals $\mu_1 - \mu_0 = 0.6 - 0.4 = 0.2$ while in the case (U) of uninformed agents the distance between the mean of the two signals is $\mu_1 - \mu_0 = 0.51 - 0.49 = 0.02$. In both cases we use a standard deviation of $\sigma_{0,1} = \sqrt{0.1}$.⁹ We conduct our simulations with $N = 100$ agents and update $T = 100$ times. To analyze the impact of the network structure on the probability of coordination on a state non-matching action, we vary the network density ρ of a random (Erdős-Rényi) graph within $\rho = [0.0, 0.95]$ in 20 steps. Each simulation is repeated $S = 1,000$ times to account for stochasticity. For all simulations we assume that the state of the world is $\theta = 0$. An overview of the parameters used can be found in Table 1.

Figure (1) shows the average final action of the system after $T = 100$ update steps for a random graph with varying densities in the informed and uninformed case for the equal weighting, neighborhood size, and relative neighborhood scenario. When the network density is very low, agents effectively act on the basis of their private signal only and the fraction of agents that choose a state non-matching action is proportional to the signal informativeness. With increasing network density, social learning sets in and the fraction of agents with a state non-matching action is reduced. With an informative signal and equal weighting there is almost no agent that chooses a state non-matching action after $T = 100$ update steps. In the neighborhood size scenario agents are more likely to follow their neighbors than in the equal weighting scenario. The initial conditions of the simu-

⁹A larger standard deviation would make the signal less informative, without changing our results qualitatively.

lation thus have a stronger effect on an agent’s decision. It will therefore take a longer time for agents to choose a state matching action. As long as the signal is not fully informative, equation (11) implies that an agent i will only choose action $x^i = 1$ if more than half of her neighbors chose that action. With increasing network density the dependence on the initial conditions will become more important, which explains the increase in the fraction of agents choosing a state non-matching action with increasing network density in the center-left panel of Figure (1). Note that this effect is not present for uninformative signals as can be seen in the right panel of Figure (1). Finally, agent i in the relative neighborhood scenario chooses $x^i = 1$ only if a relatively large fraction of her neighbors also chooses this action. The threshold for this is independent of the size of the neighborhood and only depends on the total network size.

The initial conditions become more important in the case of informative signals for the neighborhood size scenario. In order to understand how exactly our results depend on the initial conditions, Figure (2) shows the probability that a large fraction ($> 80\%$) of agents coordinate on a state non-matching action for three cases: (1) For a full sample of $S = 1,000$ simulations; (2) conditional on agents starting on average with a state matching action: $\hat{x} = \sum_i x^i(0)/N < \frac{1}{2}$; (3) conditional on agents starting on average with a state non-matching action: $\hat{x} > \frac{1}{2}$. As expected, the probability that agents coordinate on a state non-matching action drastically increases when agents initially start with a state non-matching action. However, less so in the equal weighting scenario because agents place less weight on their social beliefs and private signals are informative. A comparison of the left-center with the left-top panel in Figure (2) shows that the probability of contagion increases by a factor of roughly 200 in the neighborhood size scenario compared to the equal weighting scenario. Again, for uninformative signals this effect is not present.

While Figure (2) shows the existence of contagion even for initial actions that are state matching, the relationship between average initial and average final action is not yet quantified. Therefore, in Figure (3) we plot the average final action versus the average initial action in a density plot. We conducted a total of $S = 1,000 \times 20$ simulations and show the resulting pair of average initial and average final action as a dot with the respective coordinates. The left side of Figure (3) is for an informative signal and we draw initial actions according to the (informative) private signal. The mean of initial distributions is thus $\hat{x}_I < 0.5$, i.e. informative on average. In the right panel we show the same results for an uninformative signal and the mean of initial distributions is thus much closer to $\hat{x}_I = \frac{1}{2}$. Contagion is shown in the upper right quadrant of each subfigure. For the equal weighting scenario (top), only very few simulations yield a final average action that is state non-matching. A similar picture can be seen for the relative neighborhood (bottom) scenario. In the neighborhood size (center) scenario, however, a substantial number of simulations with an initially state matching average action yield a final state non-matching average action, confirming the existence of a contagious regime.

3 Contagious Synchronization in Endogenously Formed Networks

Banks form interbank networks endogenously. The decision whether or not two banks engage in interbank lending, e.g. in the form of agreeing on a mutual line of credit, depends in reality on many factors, including liquidity needs and counterparty risk. In the previous section we analyzed how one bank can learn about an underlying state of the world by observing the action of another bank to which it has issued an interbank credit. This additional information can create a benefit for the loan-issuing bank that constitutes a, possibly positive, externality for the lending decision. When banks coordinate on a state non-matching information, however, “learning”

about a neighboring bank's action constitutes a negative externality. The net effect of both externalities determines whether two banks are willing to engage in interbank lending. We compute the value of an additional link in three steps. First, we compute the probability that an agent chooses a state matching action, given her signal structure, private beliefs and neighbors' actions. Given this probability, we compute, second, an agent's expected utility conditional on her social belief which depends on her strategic choice to establish a link. Once an agent's expected utility with and without a link is computed, we can use the concept of pairwise stable networks to determine the equilibrium network structure.

3.1 The Probability that Agents Choose a State Matching Action

Based on the signal structure, the update of private beliefs, and neighbors' action choice we can derive the probability that an agent i chooses a state matching action. To see this, we first derive the distribution of private beliefs. The private signal structure is given as:

$$\begin{aligned} f_0(s) &= \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left(\frac{-(s - \mu_0)^2}{2\sigma_0^2}\right) \\ f_1(s) &= \frac{1}{\sqrt{2\pi}\sigma_1} \exp\left(\frac{-(s - \mu_1)^2}{2\sigma_1^2}\right) \end{aligned} \quad (17)$$

and we assume that $\sigma_0 = \sigma_1 = \sigma$. Denote the probability distribution of agent i 's private belief p^i as $f_p(p^i)$. We can then state the following:

Proposition 1 *For $\theta \in \{0, 1\}$ the probability that agent i chooses a state matching action $x^i = \theta$, given a social belief $q^i = q$ and private belief p^i , is given by:*

$$\Pr(x^i = \theta \mid q^i = q) = \int_0^{1-q} f_p(p^i \mid \theta = 0) dp^i = \int_{1-q}^1 f_p(p^i \mid \theta = 1) dp^i, \quad (18)$$

where the distribution $f_p(p^i \mid \theta = 0)$ of agent i 's private belief is given as:

$$f_p(p^i \mid \theta = 0) = \frac{\left(\left(-\frac{1-p}{p^2} - \frac{1}{p}\right) p \sigma^2\right) \exp\left(-\frac{((\mu_0 - \mu_1)^2 - 2\sigma^2 \log(\frac{1}{p} - 1))^2}{8\sigma^2(\mu_0 - \mu_1)^2}\right)}{(\sqrt{2\pi}\sigma)((1-p)(\mu_0 - \mu_1))}. \quad (19)$$

and similarly for $f_p(p^i \mid \theta = 1)$.

Proof, see Appendix (C). ■

In the absence of maturity, liquidity and counterparty risk, the value of an interbank loan, is proportional to the probability that the newly connected neighbor chooses a state-matching action and thus proportional to (18).

3.2 Agents' Expected Utility

Recall that an agent i 's utility u^i is given as:

$$u^i(x^i) = \begin{cases} 1 & \text{if } x^i = \theta \\ 0 & \text{if } x^i \neq \theta \end{cases} \quad (20)$$

The expected utility of agent i conditional on her social belief q^i is thus:

$$\bar{u}^i(x^i \mid q^i) = \Pr(x^i = \theta \mid q^i = q). \quad (21)$$

The value of a link is given by the marginal utility from establishing a link, which in turn depends on the change in the social belief q . An agent can thus influence her social belief by strategically choosing neighbors. The probability that a neighbor takes a state matching action depends in turn on the social belief that this neighbor forms about her neighbors, which leads to complex higher-order effects which we neglect in this paper. Rather, we assume that an agent i has constant beliefs about her neighbors' social beliefs q' . An agent who ignores the effect of second-nearest neighbors (i.e. neighbors of neighbors) on the social beliefs of nearest neighbors will simply assume that her neighbors' social belief is $q' = \frac{1}{2}$. This amounts to assuming

that the neighbors' actions are independent of each other. The expected utility of agent i conditional on a given q' and neighborhood K^i is given as:

$$\bar{u}^i(q', K^i) = \sum_{a \in Q^i} \Pr(q^i = a \mid q') \Pr(x^i = \theta \mid q^i = a), \quad (22)$$

where Q^i is the set of all possible values of the social belief of agent i . The first term on the right-hand side of equation (22) is the probability that agent i has a certain social belief given the social belief of her neighbors. The second term is the probability of choosing a state matching action given that social belief and given by equation (18). For a given size of the agent's neighborhood k^i , Q^i is simply $Q^i = \{n/k^i \mid n \in \mathbb{Z}, 0 \leq n \leq k^i\}$. The probability of a particular social belief can be computed by summing over the probabilities of combinations of actions chosen by the neighbors of agent i . Define the set of feasible action vectors of i 's neighbors conditional on agent i having a social belief $q^i = a$:

$$X^{ai} = \{\mathbf{x} \mid \sum_j x_j = a, x_j \in \{0, 1\}, j \in K_i\}. \quad (23)$$

i.e. X^{ai} is the set of all action vectors that are compatible with a social belief $q^i = a$. Then the probability of agent i having a private belief of $q^i = a$ given the social beliefs of all i 's neighbors, q' , is given as:

$$\Pr(q^i = a \mid q') = \sum_{\mathbf{y} \in X^{ai}} \prod_{j \in K^i} \Pr(x^j = y^j \mid q^j = q'). \quad (24)$$

Now define the probability that neighbor j chooses a state matching action as:

$$z^j = \Pr(x^j = \theta \mid q^j = q') = \int_0^{1-q^j} f_p(p^j \mid \theta = 0) dp \quad (25)$$

Note, that $f_p(p^j \mid \theta = 0)$ depends on the signal structure of neighbor j .

We can write for the probability of agent i having a private belief of $q^i = a$

in equation (24):

$$\Pr(x^j = y^j \mid q^j = q') = \begin{cases} z^j & \text{if } y^j = 0 \\ 1 - z^j & \text{if } y^j = 1 \end{cases} \quad (26)$$

If $z^j = z \forall j$ the distribution in 24 would be a simple binomial distribution. In general, however, this is not the case and we need to resort to numerical methods to compute the equilibrium network structures.

3.3 The Network Formation Process

We now have all necessary ingredients to compute the expected utility of an agent i given her neighborhood K^i and expectations about her neighbors' social belief q' . In the endogenous network formation process the agent will seek to maximize her utility by changing her neighborhood while holding q' fixed. In this section we outline an algorithm for endogenous network formation that ensures a pairwise stable network in the sense of Jackson and Wolinsky (1996):

Definition 1 *A network defined by an adjacency matrix g is called pairwise stable if*

- (i) *For all banks i and j directly connected by a link, $l^{ij} \in L$: $u^i(g) \geq u^i(g - l^{ij})$ and $u^j(g) \geq u^j(g - l^{ij})$*
- (ii) *For all banks i and j not directly connected by a link, $l^{ij} \ni L$: $u^i(g + l^{ij}) < u^i(g)$ and $u^j(g + l^{ij}) < u^j(g)$*

where the notation $g + l^{ij}$ denotes the network g with the added link l^{ij} and $g - l^{ij}$ the network with the link l^{ij} removed. When maintaining a link is costly, there will be some network density that depends on the cost $c > 0$ per link. The marginal utility of an additional link decreases with the number of links because the expected utility is bounded by 1 (the pay-off is 1 and the probability of choosing the correct action is less than, or equal to, 1).

The algorithm to ensure a pairwise stable equilibrium starts by choosing a random agent i from the set of agents N . Then, choose a second agent j from the set of agents $N \setminus K^i$ that are not yet neighbors of i . Agents are chosen with the following probability:

$$w^j = \exp(\beta E^j)/Z, \quad (27)$$

where $Z = \sum_k w^k$ is a normalization constant and $E^j = |\frac{1}{2} - \mu_0^j|$ is a proxy for agent j 's signal strength. For $\beta = 0$ i chooses the new agent with equal probability. While this makes it more likely that agent i considers forming a link with agent j when j has a higher signal strength, it does not imply that such a link is actually formed. This decision is solely based on the utility that both i and j obtain from establishing the link.

Now, let $K'^i = K^i \cup j$, i.e. the neighborhood of agent i after adding agent j , and similarly $K'^j = K^j \cup i$. The marginal utilities of adding j and i to the respective neighborhoods are then:

$$\begin{aligned} \Delta \bar{u}^i(q', K'^i, K^i) &= \bar{u}^i(q', K'^i) - \bar{u}^i(q', K^i) \\ \Delta \bar{u}^j(q', K'^j, K^j) &= \bar{u}^j(q', K'^j) - \bar{u}^j(q', K^j) \end{aligned} \quad (28)$$

Given the marginal utilities of agents i and j and their cost of maintaining link c^i and c^j the agents will form a link if $\Delta \bar{u}^i(q', K'^i, K^i) > c^i$ and $\Delta \bar{u}^j(q', K'^j, K^j) > c^j$. If $\Delta \bar{u}^i(q', K'^i, K^i) > c^i$ and $\Delta \bar{u}^j(q', K'^j, K^j) < c^j$, the algorithm selects the least informative agent in the neighborhood of j :

$$l = \underset{m \in K^j}{\operatorname{argmin}} E^m. \quad (29)$$

Now, define $K''^j = K'^j \setminus l$. If $\Delta \bar{u}^i(q', K'^i, K^i) > c^i$ and $\Delta \bar{u}^j(q', K''^j, K^j) > c^j$, form the link l^{ij} and remove the link l^{jk} (and similarly for $i \rightarrow j$ and $j \rightarrow i$). Otherwise, don't form the link. If $\Delta \bar{u}^i(q', K'^i, K^i) < c^i$ and $\Delta \bar{u}^j(q', K'^j, K^j) < c^j$ repeat the previous step, i.e. consider removing the

least informative neighbor and re-evaluate the utilities.

3.4 Equilibrium Networks

The endogenously formed network in an economy with identically informed agents and positive cost c of maintaining a link is a simple Erdős-Rényi network with a network density depending on the signal structure and link cost. To analyze more realistic situations, we can harvest the strengths of multi-agent simulations. In the following we therefore assume that agents are heterogeneously informed about the underlying state of the world: a few “informed” agents have relatively precise signals and low costs of maintaining a link, while many “uninformed” agents have relatively imprecise signals and higher cost of maintaining a link. Table 2 summarizes the parameters we are using for the rest of this section.

In order to compare the dynamics on the endogenous networks to the ER networks we run the following simulations. We first create 1,000 networks using the network formation algorithm described above. Then, we run the social learning algorithm described in Section 2 while holding the network structure constant throughout. The underlying assumption is that banks are updating their investment decisions faster than the network structure changes. This can be empirically corroborated by looking at the term structure of interbank lending. While 90% of the turnover in interbank markets is overnight, about 90% of exposures between banks stems from the term segments.

We also run the social learning algorithm with an initialization bias in which we set the initial action of all agents to some pre-defined value. To assess the efficiency of the endogenous network formation, we compare the performance of the endogenously formed networks to the performance of Erdős-Rényi networks. All simulations in this section are conducted using the equal weighting scenario. An example of a resulting network structure can

be found in Figure 5 and the degree distribution of the endogenously formed networks in 1,000 is shown in Figure 6. The degree distribution of the endogenously formed networks is clearly bimodal. One peak corresponds to the uninformed nodes with small degree while the second peak corresponds to the informed nodes with high degree. Figure 7 shows the distribution of the final action for the 1,000 simulations conducted. A clear improvement over the Erdős-Rényi networks can be seen, highlighting the importance of core banks with more precise private signals. Since core banks are highly interconnected, there is a higher chance that they are in the neighborhood of a peripheral bank (as opposed to a peripheral bank being in the neighborhood of another peripheral bank) which increases the precision of peripheral banks' social belief.

To further understand the difference between endogenously formed and random networks, we analyze the time it takes learning to converge. We assume the learning has converged at time t if:

$$\sum_i |x_i(t) - x_i(t + \Delta t)|/N < 0.05, \text{ where } \Delta t = 15. \quad (30)$$

It can be seen from Figure (8) that, except for very long convergence times, the system always converges faster in the endogenous network case than in the Erdős-Rényi case. Note, that this simulation was conducted without initialization bias, i.e. with average initial action of $\frac{1}{2}$. Finally, the probability of contagion as a function of an initialization bias, i.e. as a function of average initial action is shown in Figure (9). Again, the picture is unanimously showing that the probability of contagion, i.e. the probability that more than 80% of agents coordinate on a state non-matching action is significantly smaller in the endogenous network case than in the case of a random Erdős-Rényi graph.

4 Conclusion

This paper develops a model of contagious synchronization of bank's investment strategies. Banks are connected via mutual lines of credit and endogenously choose an optimal network structure. They receive a private signal about the state of the world and observe the strategies of their counterparties. When banks observe the actions of more peers they put more weight on their social belief. We compare three scenarios of weighting functions. First, in the equal weighting scenario, agents place equal weights on their private and social belief. Second, in the neighborhood scenario agents place proportionately more weight on the social signal when the size of the neighborhood increases. Third, in the relative neighborhood scenario agents place more weight on the social belief if their neighborhood constitutes a larger fraction of the overall network. Social learning increases the probability of choosing a state matching action and thus agents' utility. When agents strategically choose their neighbors they take the additional utility from learning into account. The more neighbors a given agent has, the lower is the marginal utility from another link and the network endogenously reaches an equilibrium configuration.

We obtain two results which are policy relevant. First, in a complex financial system where agents cannot take the action of all their peers into account when taking an investment decision, the probability of contagious synchronization depends on two things: (i) the weighting between the private and social belief; and (ii) the density of the financial network. Our model thus relates two empirically relevant sources of systemic risk: common shocks interbank market freezes. Second, the probability of contagious synchronization is substantially reduced when agents internalize the positive effects of social learning in a strategic decision with whom to form a link. The benefit from learning is reduced when the private signals about the state of the world are less informative.

The model has a number of interesting extensions. One example is the case with two different regions that can feature differing states of the world. Such an application could capture a situation in which banks in two countries (one in a boom, the other in a bust) can engage in interbank lending within the country and across borders. This would provide an interesting model for the current situation within the Eurozone. The model so far features social learning but not individual learning. Another possible extension would be to introduce individual learning and characterize the conditions under which the contagious regime exists. Finally, the model can be applied to real-world interbank network and balance sheet data to test for the interplay of contagious synchronization and endogenous network structure.

One drawback of the model is that there is no closed-form analytical solution for the benefit a bank obtains through learning from a peer that takes into account higher order effects. This benefit will depend on whether or not a neighboring bank chose a state matching on state non-matching action in the previous period and thus on the social belief of neighboring banks. In the former case, the benefit will be positive, while in the latter case it will be negative. Agents have ex ante no way of knowing what action a neighboring bank selected until the state of the world is revealed ex post. Finding such a closed-form solution is beyond the scope of the present paper which focuses on the application in an agent-based model, but would provide a fruitful exercise for future research.

References

- Acemoglu, D., Bimpikis, K., Ozdaglar, A., 2014. Dynamics of information exchange in endogenous social networks. *Theoretical Economics* 9, 41 – 97.
- Acemoglu, D., Dahleh, M.A., Lobel, I., Ozdaglar, A., 2011. Bayesian learning in social networks. *Review of Economic Studies* 78, 1201–1236.
- Alevy, J.E., Haigh, M.S., List, J., 2007. Information cascades: Evidence from an experiment with financial market professionals. *The Journal of Finance* 67, 151–180.
- Allen, F., Gale, D., 2000. Financial contagion. *Journal of Political Economy* 108, 1–33.
- Bala, V., Goyal, S., 1998. Learning from neighbors. *The Review of Economic Studies* 65, 559–621.
- Bala, V., Goyal, S., 2000. A noncooperative model of network formation. *Econometrica* 68, 1181–1231.
- Bandt, O., Hartmann, P., Peydró, J.L., 2009. Systemic Risk in Banking: An Update. *The Oxford Handbook of Banking* (Oxford University Press, UK) .
- Banerjee, A., 1992. A simple model of herd behaviour. *The Quarterly Journal of Economics* 107, 797–817.
- Battiston, S., Gatti, D.D., Gallegati, M., Greenwald, B.C., Stiglitz, J., 2012. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control* 36, 1121–1141.
- Bernhardt, D., Campello, M., Kutsoati, E., 2006. Who herds? *Journal of Financial Economics* 80, 657–675.

- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100, 992–1026.
- Bluhm, M., Faia, E., Krahnen, J.P., 2013. Endogenous Banks’ Networks, Cascades, and Systemic Risk. SAFE Working Paper 12. Goethe University Frankfurt.
- Caballero, R., 2012. A Macro Tragedy Waiting to Happen (in a VaR sense). Keynote Speech, TCMB Conference 2012. Türkiye Cumhuriyet Merkez Bankası.
- Castiglionesi, F., Navarro, N., 2007. Fragile Financial Networks. CentER Discussion Paper No. 2007-100. Tilburg University.
- Chang, E.C., Cheng, J.W., Khorana, A., 2000. An examination of herd behavior in equity markets: An international perspective. *Journal of Banking and Finance* 24, 1651–1679.
- Chiang, T.C., Zheng, D., 2010. An empirical analysis of herd behavior in global stock markets. *Journal of Banking and Finance* 34, 1911–1921.
- Cipriani, M., Guarino, A., 2014. Estimating a structural model of herd behavior in financial markets. *The American Economic Review* 104, 224–251.
- Craig, B., von Peter, G., 2014. Interbank tiering and money center banks. *Journal of Financial Intermediation* In press.
- DeGroot, M., 1974. Reaching a consensus. *Journal of the American Statistical Association* 69, 118–121.
- DeMarzo, P., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics* 118, 909–968.

- Freixas, X., Parigi, B.M., Rochet, J.C., 2000. Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of Money, Credit and Banking* 32, 611–38.
- Gale, D., Kariv, S., 2003. Bayesian learning in social networks. *Games and Economic Behaviour* 45, 329–346.
- Georg, C.P., 2013. The effect of the interbank network structure on contagion and common shocks. *Journal of Banking and Finance* 37, 2216–2228.
- Haldane, A.G., Madouros, V., 2012. The Dog and the Frisbee. Speech. Federal Reserve Bank of Kansas City’s 36th Economic Policy Symposium, *"The Changing Policy Landscape"*, Jackson Hole, Wyoming.
- Iori, G., Jafarey, S., Padilla, F., 2006. Systemic risk on the interbank market. *Journal of Economic Behavior & Organization* 61, 525–542.
- Jackson, M.O., Wolinsky, A., 1996. A strategic model of social and economic networks. *Journal of Economic Theory* 71, 44–74.
- Ladley, D., 2013. Contagion and risk sharing on the interbank market. *Journal of Economic Dynamics and Control* 37, 1384–1400.
- Nier, E., Yang, J., Yorulmazer, T., Alentorn, A., 2007. Network models and financial stability. *Journal of Economic Dynamics and Control* 31, 2033–2060.
- Poledna, S., Thurner, S., Farmer, J.D., Geanakoplos, J., 2014. Leverage-induced systemic risk under basle ii and other credit risk policies. *Journal of Banking and Finance* 42, 199–212.

A Tables

Variable	Description	Value		
		1 - (I)	2 - (U)	3 - (H)
N	Number of agents	100	100	100
μ_0	Average signal for $\theta = 0$	0.4	0.49	0.3
μ_1	Average signal for $\theta = 1$	0.6	0.51	0.7
σ_0	Standard deviation of signal for $\theta = 0$	$\sqrt{0.1}$	$\sqrt{0.1}$	$\sqrt{0.1}$
σ_1	Standard deviation of signal for $\theta = 1$	$\sqrt{0.1}$	$\sqrt{0.1}$	$\sqrt{0.1}$
T	Number of iterations of updater	100	100	100
ρ	Density of ER network	$[0, 0.95]$	$[0, 0.95]$	0.5
p	Probability of being informed	NA	NA	$[0.1, 0.9]$
S	# of simulations per param. config.	1000	1000	1000

Table 1: Parameters for runs for the case of informed (I), uninformed (U), and heterogeneous (H) agents.

Notation	Description	Value
N	Number of agents	30
N_I	Number of informed agents	4
N_U	Number of uninformed agents	26
μ_{0I}	Average signal for $\theta = 0$ for informed agents	0.3
μ_{1I}	Average signal for $\theta = 1$ for informed agents	0.7
σ_{0I}	Standard deviation of signal for $\theta = 0$ for informed agents	$\sqrt{0.1}$
σ_{1I}	Standard deviation of signal for $\theta = 1$ for informed agents	$\sqrt{0.1}$
μ_{0U}	Average signal for $\theta = 0$ for uninformed agents	0.4
μ_{1U}	Average signal for $\theta = 1$ for uninformed agents	0.7
σ_{0U}	Standard deviation of signal for $\theta = 0$ for uninformed agents	$\sqrt{0.1}$
σ_{1U}	Standard deviation of signal for $\theta = 1$ for uninformed agents	$\sqrt{0.1}$
c_I	Cost per link for informed agent	0
c_U	Cost per link for informed agent	0.1
T_C	Number of iterations of network algorithm	400
q'	Agent's belief of average action of neighbors of neighbors	0.5
β	Intensity of choice in agent selection	30
T	Number of iterations of action updater	100
n	Number of endogenously formed networks used for simulation	1000
S	Number of simulations per parameter configuration	100
ρ	Average density of ER networks	0.08

Table 2: Parameters for network formation and runs with endogenous networks.

B Figures

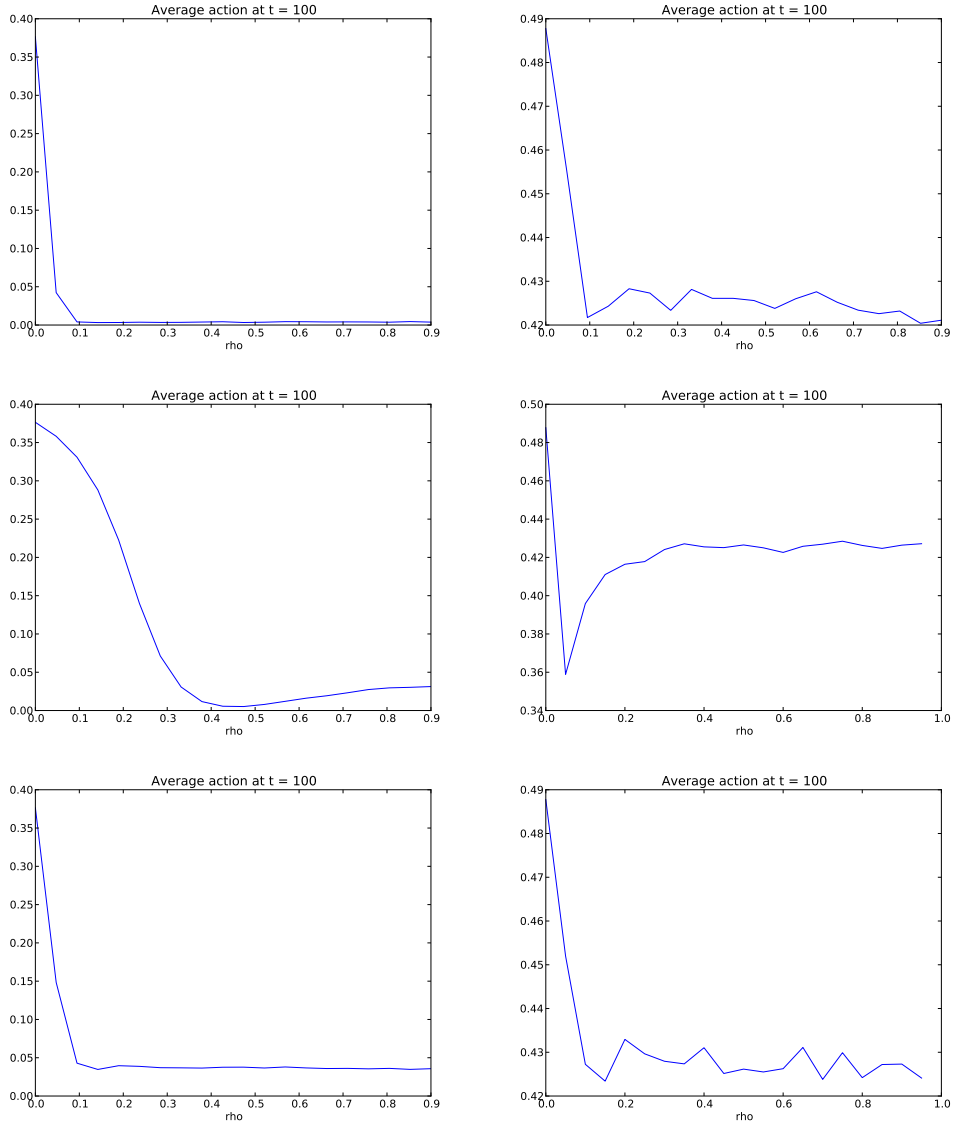


Figure 1: Average final action of agents as a function of the network density ρ of a random graph for the informed (left) and uninformed (right) case. Top: Equal weighting scenario; Center: Neighborhood size scenario; Bottom: relative neighborhood scenario.

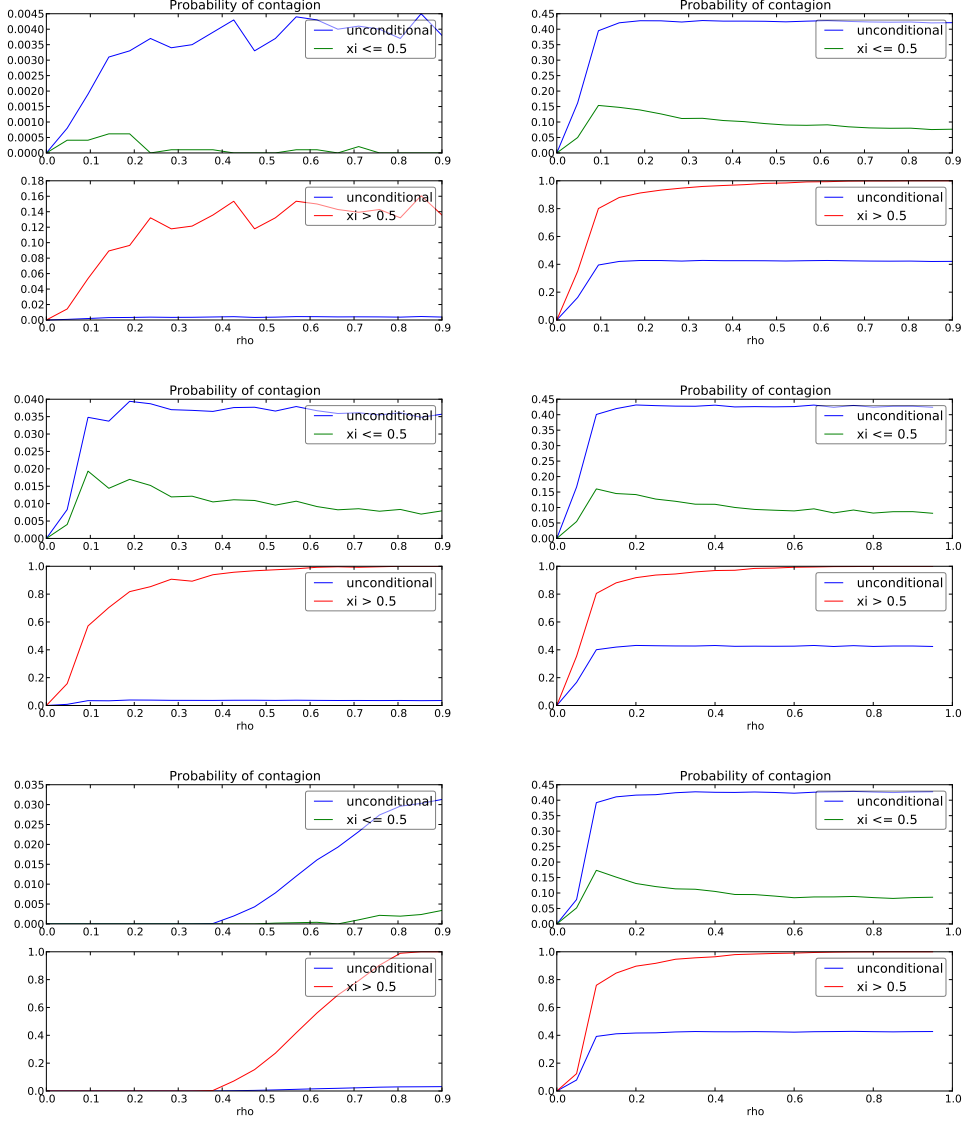


Figure 2: Fraction of simulations per parameter configuration S (1000) in which agents synchronize on the state non matching action (more than 80% of agents choose state non-matching action) as a function of network density ρ of a random graph for the informed (left) and uninformed (right) case. Top: Equal weighting scenario; Center: Neighborhood size scenario; Bottom: relative neighborhood scenario. We distinguish three cases: (1) unconditional: we compute the fraction based on the full sample S . (2) conditional $\hat{x} \leq \frac{1}{2}$: we compute the fraction based on the sub-set of simulations in which the average initial action $\hat{x} = \sum_i x^i(0)/N \leq \frac{1}{2}$, i.e. when the agents start with a state matching action. (3) conditional $\hat{x} > \frac{1}{2}$: we compute the fraction based on the sub-set of simulations when the agents start with a state non matching action.

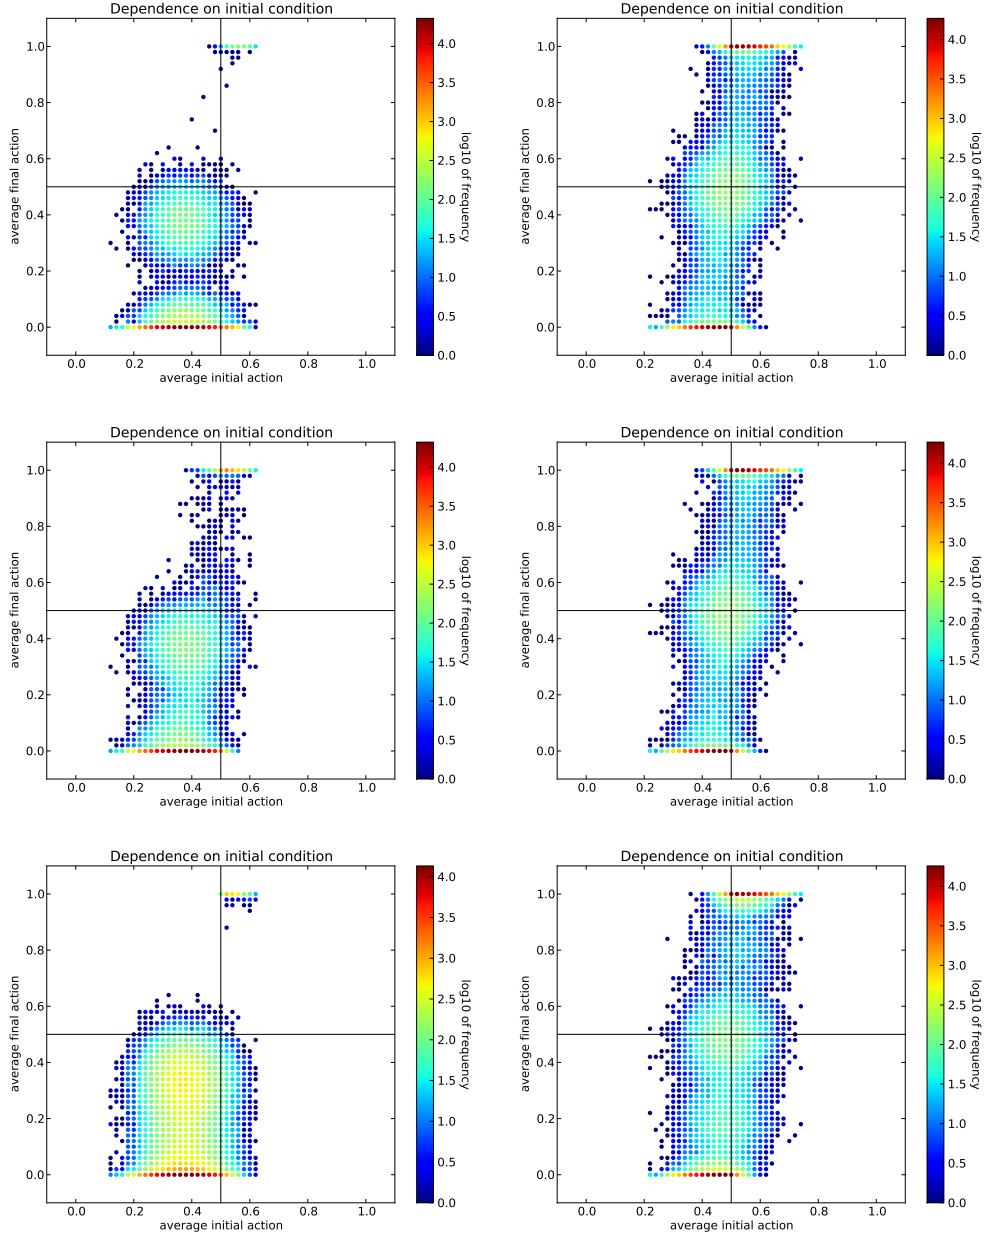


Figure 3: Average final action $\hat{x}_F = \sum_i x^i(T)/N$ versus the average initial action $\hat{x}_I = \sum_i x^i(0)/N$ for the informed (left) and uninformed (right) case. Top: Equal weighting scenario; Center: Neighborhood size scenario; Bottom: relative neighborhood scenario. Data points are averages over $S = 1000$ simulations and all network densities ρ (20 values equally distributed over the interval $[0, 0.95]$). The color code indicates the frequency with which a point occurs in the sample (total size 20×1000), the scale of the color code is logarithmic of base 10.

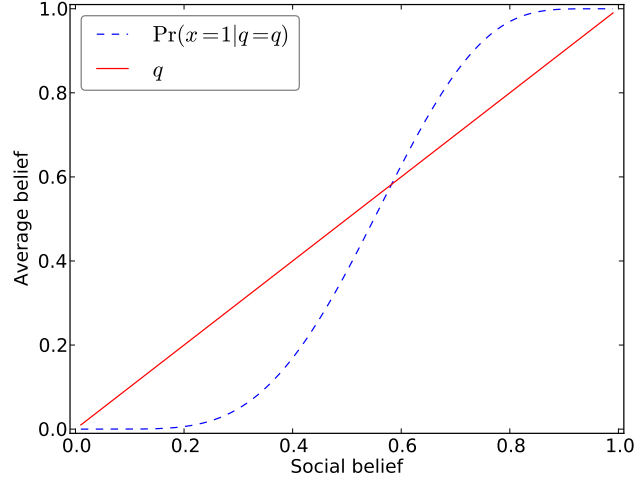


Figure 4: Probability of choosing state matching action (equivalent to average social belief and average action in mean field) given that the state of the world is $\theta = 0$, $f_p(p^i | \theta = 0)$ for $\mu_0 = 0.4$. Furthermore we use $\mu_1 = 1 - \mu_0$ and $\sigma^2 = 0.1$. The intersections between the probability function and the diagonal mark the fixed points of the dynamical system.

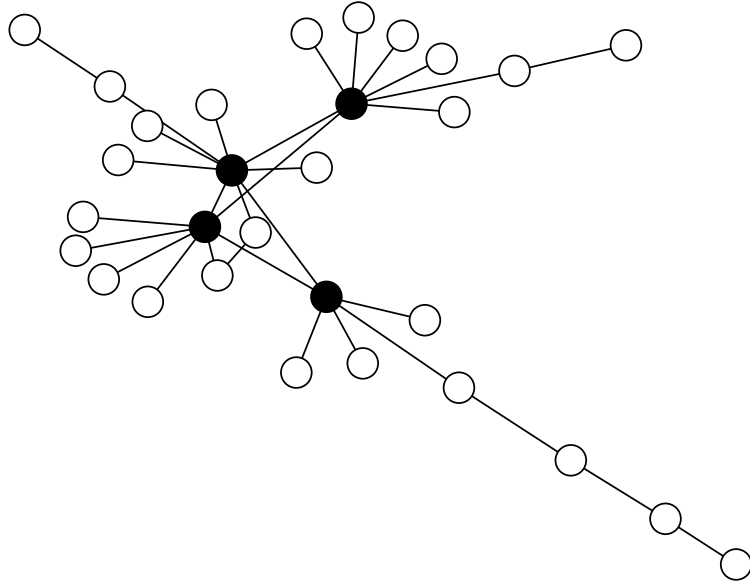


Figure 5: Example network from ensemble of $n = 1000$ endogenously formed networks. Black nodes are “informed” agents, while white nodes are “uninformed”.

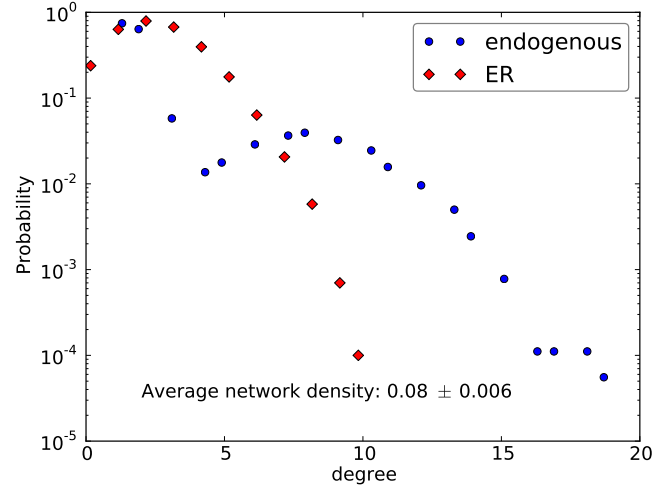


Figure 6: Degree distribution of ensemble of $n = 1000$ endogenously formed networks compared to ER networks with same average density.

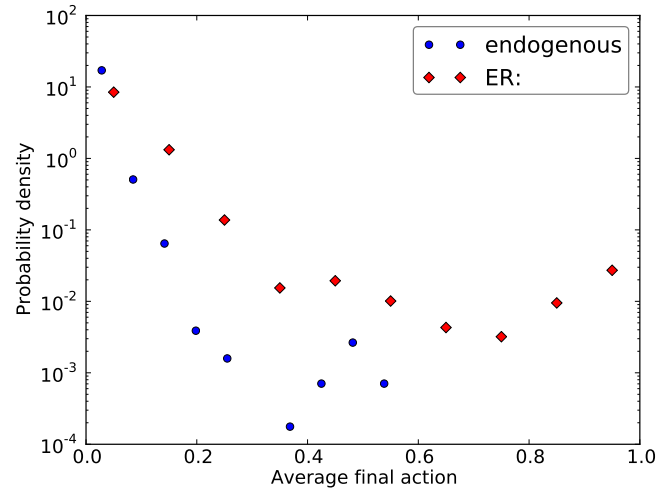


Figure 7: Distribution of final action in ER networks vs. endogenous networks. This is without bias, i.e. the initial action is random based on the private belief only.

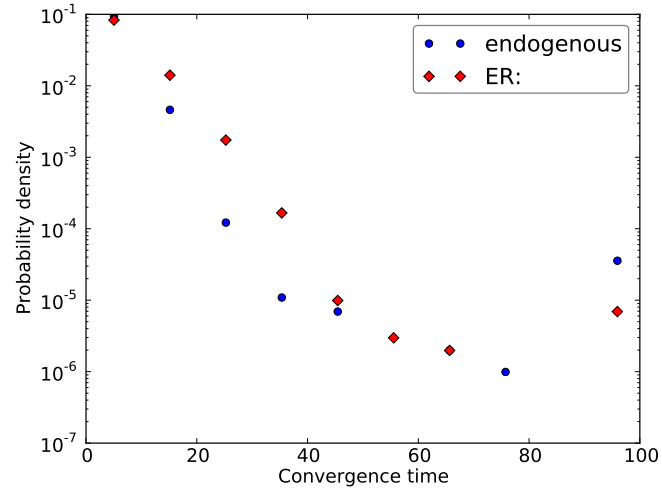


Figure 8: Distribution of convergence time ER networks vs. endogenous networks. This is without bias, i.e. the initial action is random based on the private belief only.

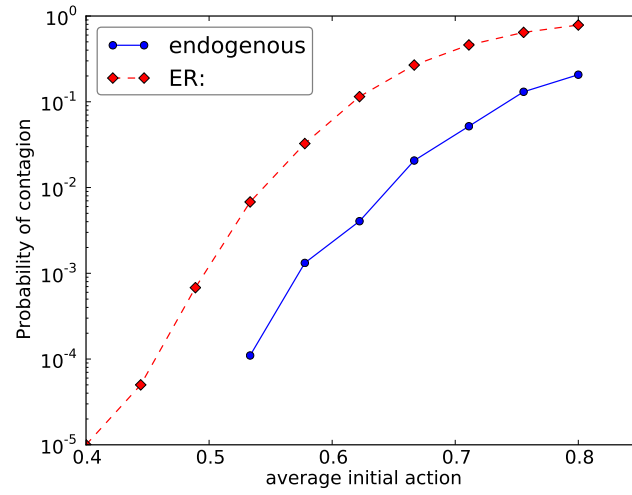


Figure 9: Probability of contagion vs. initialization bias. Missing values correspond to zero frequency.

C Proofs

Proof of Proposition (1). We use the notation $f_P(p^i \mid \theta = 0)$ to indicate that the functional form of the probability distribution of the private belief p^i has been derived assuming that $\theta = 0$. Now, let $s(p^i)$ be the inverse of the private belief:

$$s(p^i) = \frac{\mu_i^2 - \mu_1^2 + 2\sigma^2 \log\left(\frac{1-p^i}{p^i}\right)}{2(\mu_0 - \mu_1)} \quad (31)$$

The distribution of the private belief can be computed as follows:

$$f_p(p^i) = \frac{\partial s(p^i)}{\partial p^i} f_s(s(p^i)). \quad (32)$$

where the probability density function for signal s is given as:

$$f_s(s) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(s - \mu_0)^2}{2\sigma^2}\right) \quad (33)$$

and the private belief $p^i(s)$ is given by Equation (8). Substituting in the expression for $s(p^i)$ and computing the partial derivative we obtain:

$$f_p(p^i \mid \theta = 0) = \frac{\left(\left(\frac{1-p^i}{p^i} - 1\right) \sigma^2\right) \exp\left(-\frac{\left((\mu_0 - \mu_1)^2 - 2\sigma^2 \log\left(\frac{1}{p^i} - 1\right)\right)^2}{8\sigma^2(\mu_0 - \mu_1)^2}\right)}{(\sqrt{2\pi}\sigma) ((1 - p^i)(\mu_0 - \mu_1))} \quad (34)$$

Example distributions of the private belief are shown in Figure (10). We have $\mu_1 = 1 - \mu_0$ and $\sigma^2 = 0.1$. Note, that the majority of the probability density of the private belief is to the left of 0.5 in all cases. Therefore, the private signal tends to produce private beliefs that yield the state matching action. If we increase the $|\mu_0 - 0.5|$ the distribution becomes more skewed towards the actual state of the world. Hence the private belief becomes more informative.

Now that we have defined the pdf of the private belief we can compute the

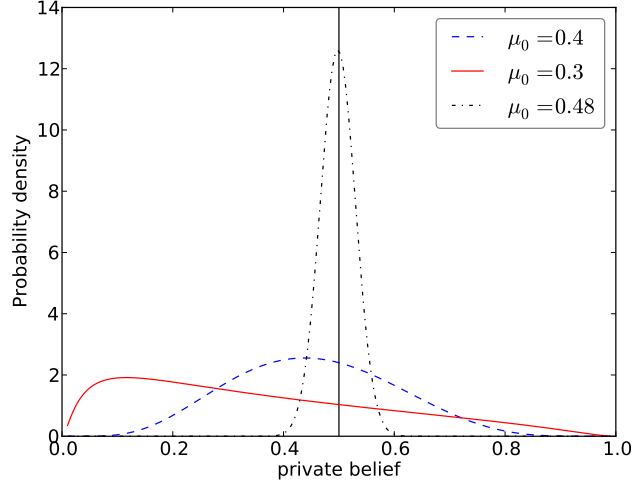


Figure 10: Probability density function of the private belief given that the state of the world is $\theta = 0$ $f_P(p \mid \theta = 0)$ for three values of $\mu_0 \in \{0.3, 0.4, 0.48\}$.

probability that the agent chooses $x^i = 0$ given some social belief $q = q^i$ as:

$$\Pr(x^i = 0 \mid q^i = q) = \int_0^{1-q} f_p(p^i \mid \theta = 0) dp^i, \quad (35)$$

where we use the notation $f_p(p^i \mid \theta = 0)$ to indicate that the functional form of f_p has been derived assuming that $\theta = 0$. This result generalizes to all θ due to the symmetry of the signal structure. It can be shown that:

$$\Pr(x^i = \theta \mid q^i = q) = \int_0^{1-q} f_p(p^i \mid \theta = 0) dp^i = \int_{1-q}^1 f_p(p^i \mid \theta = 1) dp^i. \quad (36)$$

Therefore we have derived an expression for the probability of choosing the correct action that does not depend on the actual state of the world but only on the signal structure and the social belief.