

Georg, Co-Pierre

Working Paper

Contagious herding and endogenous network formation in financial networks

ECB Working Paper, No. 1700

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Georg, Co-Pierre (2014) : Contagious herding and endogenous network formation in financial networks, ECB Working Paper, No. 1700, ISBN 978-92-899-1108-5, European Central Bank (ECB), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/154133>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM



In 2014 all ECB publications feature a motif taken from the €20 banknote.



WORKING PAPER SERIES

NO 1700 / JULY 2014

CONTAGIOUS HERDING AND ENDOGENOUS NETWORK FORMATION IN FINANCIAL NETWORKS

Co-Pierre Georg

ECB LAMFALUSSY FELLOW-
SHIP PROGRAMME

NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Lamfalussy Fellowships

This paper has been produced under the ECB Lamfalussy Fellowship programme. This programme was launched in 2003 in the context of the ECB-CFS Research Network on “Capital Markets and Financial Integration in Europe”. It aims at stimulating high-quality research on the structure, integration and performance of the European financial system.

The Fellowship programme is named after Baron Alexandre Lamfalussy, the first President of the European Monetary Institute. Mr Lamfalussy is one of the leading central bankers of his time and one of the main supporters of a single capital market within the European Union.

Each year the programme sponsors five young scholars conducting a research project in the priority areas of the Network. The Lamfalussy Fellows and their projects are chosen by a selection committee composed of Eurosystem experts and academic scholars. Further information about the Network can be found at <http://www.eufinancial-system.org> and about the Fellowship programme under the menu point “fellowships”.

Acknowledgements

I would like to thank Toni Ahnert, Jean-Edouard Colliard, Jens Krause, Tarik Roukny, an anonymous referee, 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. Outstanding research assistance by Christoph Ay-manns is gratefully acknowledged. This paper has been prepared by the author under the Lamfalussy Fellowship Program sponsored by the ECB. The views expressed in this paper do not necessarily reflect the views of Deutsche Bundesbank or the ECB.

Co-Pierre Georg

Deutsche Bundesbank and University of Cape Town; e-mail: co-pierre.georg@bundesbank.de.

© European Central Bank, 2014

Address	Kaiserstrasse 29, 60311 Frankfurt am Main, Germany
Postal address	Postfach 16 03 19, 60066 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Internet	http://www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors. This paper can be downloaded without charge from <http://www.ecb.europa.eu> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=2464109. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB’s website, <http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html>

ISSN	1725-2806 (online)
ISBN	978-92-899-1108-5 (online)
EU Catalogue No	QB-AR-14-074-EN-N (online)

Abstract

When banks choose similar investment strategies, the financial system becomes vulnerable to common shocks. Banks decide about their investment strategy ex-ante based on a private belief about the state of the world and a social belief formed from observing the actions of peers. When the social belief is strong and the financial network is fragmented, banks follow their peers and their investment strategies synchronize. This effect is stronger for less informative private signals. For endogenously formed interbank networks, however, less informative signals lead to higher network density and less synchronization. It is shown that the former effect dominates the latter.

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

JEL Classification: G21, C73, D53, D85

Non-technical Summary

In order to derive optimal investment strategies, banks gather information from a private signal and a social signal obtained by observing peers. With all this information available, how could it be that so many banks chose an investment strategy which was at odds with the current state of the world in the run-up to the financial crisis? This paper develops a simple model of banking behavior in which banks ex-ante coordinate their investment strategy on a joint action that is not necessarily matching the state of the world. The model exhibits two states of the world and banks choose one of two investment strategies. They receive a positive utility if their action (i.e. their investment strategy) matches the state of the world and nothing otherwise. Banks receive both private signals about the state of the world and observe the actions of other banks to which they are connected via mutual lines of credit.

I obtain three results. First, with an exogenously given network structure, I show the existence of a *contagious regime* in which all banks ex-ante choose an investment strategy that makes them susceptible to an ex-post common shock. Contagious synchronization on a state non-matching action exists for interim levels of interconnectedness only, which relates two distinct sources of systemic risk: common shocks and interbank market freezes. When there is heightened ex-ante uncertainty about the state of the world there exists a higher probability that the strategies of all banks become synchronized, making them vulnerable to a common shock. This effect is larger when banks are less interconnected, i.e. when interbank markets dry up. I analyze the impact of different network topologies on the existence of the contagious regime, showing that the existence of the contagious regime depends on the properties of the network irrespective of the network type. I also show that a contagious regime occurs even when some agents are significantly better informed than others.

Second, I characterize the pairwise stable equilibrium network structures when banks endogenously choose their counterparties and receive a highly informative signal. When banks receive more informative signals about the underlying state of the world they put less value on liquidity coinsurance and more on the threat of ex-post contagion via counterparty risk. The resulting network density decreases with the informativeness of the signal structure. The model presented in this paper thus provides an additional explanation for the persistence of interbank markets in the face of heightened uncertainty.

Third, I investigate the full model with social learning and endogenous network formation and show that the size of the contagious regime is reduced with increasing signal informativeness. This implies that the size of the contagious regime is more reduced through increased informativeness than increased through the reduction in network density that follows from increasing informativeness.

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 ex-post common shocks.¹ A case at hand is the financial crisis of 2007/2008 when many banks invested in mortgage backed securities in anticipation that the underlying mortgages, many of which being US subprime mortgages, would not simultaneously depreciate in value. Fatally, this assumption turned out to be incorrect, and systemic risk ensued. How could so many banks choose the “wrong”, i.e. non-optimal given the state of the world, investment strategy, although they carefully monitor both economic fundamentals and the actions of other banks?

This paper presents a model in which financial intermediaries herd ex-ante and synchronize their investment strategy on a state non-matching strategy 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 every point in time with a certain probability p . 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 strategies of their peers (but not of other banks). Based on both observations, they form a belief about the state of the world and choose their action accordingly.

The model presented in this paper is in essence a simple model of Bayesian learning in social networks but deviates from the existing literature (e.g.

¹Common shocks are a form of systemic risk in the broad sense of See Bandt et al. (2009).

Gale and Kariv (2003), Acemoglu et al. (2011)) along two dimensions. First, instead of observing the actions of one peer at a time, each time adjusting their strategy accordingly, I assume that agents average over their peers' actions in the previous period.² 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.³ Second, I allow banks to endogenously choose their set of peers in an extension of the baseline model. Here the underlying assumption is that banks have limited resources and do not monitor the actions of all other banks, but only from a strategically chosen subset. They receive utility from being interconnected through a learning effect. Banks trade-off benefits from this coinsurance with a potential counterparty risk from peers choosing a state non-matching action (i.e. ex-post interbank contagion), being short on liquidity to rebalance their portfolio and thus drawing on the credit line.⁴ Finally, banks suffer larger losses when they chose a state non-matching action and the financial network is more densely interconnected through an amplification effect occurring e.g. when many agents rebalance their portfolios simultaneously, thereby triggering a fire-sale. The resulting endogenous network structure is pairwise stable in the sense of Jackson and Wollinsky (1996). The model is implemented as an agent-based model (ABM) of the financial system.

I obtain three results. First, in the limiting case with a fixed and exogenously given network structure, I show the existence of a *contagious regime* in which all banks ex-ante coordinate on a state non-matching action (i.e. they choose an investment strategy that performs badly when the state of

²Contagious synchronization on a state non-matching action occurs before the state of the world is revealed and is thus an ex-ante form of contagion.

³This assumption renders the agents boundedly rational, albeit mildly so, as for example DeMarzo et al. (2003) argue.

⁴This simple setting introduces counterparty risk since a bank that chose a state non-matching action draws on a mutual line of credit, reducing liquidity available at the peer which increases the potential for a liquidity shortage if the peer chose a state non-matching action as well.

the world is revealed). In a fully connected network there is no contagious synchronization on a state non-matching action, since signals are informative and every bank observes the actions of every other bank, which on average will be state matching. In a network that is not fully connected there is a chance that a bank i is connected to a set of peers K^i that chose a state-non-matching action on average. As there is only social learning in the model, the social belief of bank i might exceed the private belief and bank i chooses a state non-matching action in the next time step. This increases the chance that another bank now has a neighborhood in which a majority of banks chose a state-non-matching action, and the process is repeated until finally all banks chose a state-non-matching action.⁵ This iterative process is the driving force of bank behaviour in the contagious regime. The probability of entering a contagious regime is smaller for networks that are more connected.

This result is of particular interest for policy makers as it relates two sources of (ex-post) systemic risk: common shocks and interbank market freezes. When there is heightened ex-ante uncertainty about the state of the world there exists a higher probability that the strategies of all banks become synchronized, making them vulnerable to a common shock. This effect is larger when banks are less interconnected, i.e. when interbank markets dry up. Caballero (2012) documents a higher correlation amongst various asset classes in the world in the aftermath of the Lehman insolvency, i.e. during times of heightened uncertainty. This can be understood by a synchronization of bank's investment strategies for which the model provides a simple rationale.

Two extensions of the main result are discussed. First, I analyze the impact of different network topologies on the existence of the contagious regime,

⁵Such *informational cascades* are a well-documented empirical phenomenon. See, for example, Alevy et al. (2007), Bernhardt et al. (2006), Chang et al. (2000), and Chiang and Zheng (2010).

showing that the existence of the contagious regime depends on the properties of the network (i.e. shortest average path-length) rather than on the type (i.e. random, scale-free, small-world). Second, I show that a contagious regime occurs even when some agents are significantly better informed than others.

Second, turning to the extension of endogenously chosen networks, when banks receive highly informative signals I characterize the pairwise stable equilibrium network structures. When banks receive more informative signals about the underlying state of the world, they put less value on liquidity coinsurance and more on the threat of contagion through counterparty risk. The resulting network density decreases with the informativeness of the signal structure. This result is of particular interest when applied to the financial crisis of 2007/2008. Although there is a substantial body of theoretical literature on interbank market freezes (see e.g. Acharya and Skeie (2011), Gale and Yorulmazer (2013), Acharya et al. (2011)), the empirical evidence is mixed. Acharya and Merrouche (2013) provide evidence of liquidity hoarding by large settlement banks in the UK on the day after the Lehman insolvency. Afonso et al. (2011), however, analyze the US overnight interbank market and show that while the market was stressed in the aftermath of the Lehman insolvency (i.e. loan terms become more sensitive to borrower characteristics), it did not freeze. Gabrieli and Georg (2013) obtain similar results for the euroarea, showing that interbank markets did not freeze, but banks were shortening the maturity of their interbank exposures. The model presented in this paper provides an additional explanation for the persistence of interbank markets in the face of heightened uncertainty.

Third, I investigate the full model with social learning and endogenous network formation and show that the size of the contagious regime is reduced with increasing signal informativeness. This implies that the size of the contagious regime is more reduced through increased informativeness than

increased through the reduction in network density that follows from increasing informativeness.

This paper relates to three strands of literatures. First, and foremost, the paper develops a financial multi-agent simulation in which agents learn through endogenously formed interconnections. 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.⁶ Slight deviations from these models can be found, for example, in Bluhm et al. (2013), Ladley (2013), and Georg (2013) who employ different equilibrium concepts. While Bluhm et al. (2013) uses a tatonnement process to obtain an equilibrium interest rate on the interbank market, Georg (2013) uses a rationing model and introduces a central bank to enable rationed banks to access central bank facilities. Ladley (2013) uses a genetic algorithm to find equilibrium values for bank balance sheets. 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 Wollinsky (1996) and Bala and Goyal (2000). The paper in this literature closest to mine is Castiglionesi and Navarro (2011). The authors study the formation of endogenous networks in a banking network with microfounded banking behaviour. Unlike Castiglionesi and Navarro (2011), however, my model uses a 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

⁶Closely 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 and Allen and Babus (2009) for a more recent survey. The vast majority of models in this literature considers a fixed network structure only (see, amongst various others, Gai and Kapadia (2010), Gai et al. (2011), Battiston et al. (2012), Haldane and May (2011)).

not present in the work of Castiglionesi and Navarro (2011). Other papers in this literature include Babus (2011), Castiglionesi and Wagner (2013), Babus and Kondor (2013), and Cohen-Cole et al. (2012).

Finally, the Bayesian learning part of the model is closely related to the literature on Bayesian learning in social networks. Acemoglu et al. (2011) study a model of sequential learning in a social network where each agent receives a private signal about the state of the world and observe past actions of their neighbors. Acemoglu et al. (2011) show that asymptotic learning (i.e. choosing the state-matching action with probability 1) occurs when private beliefs are unbounded. My model, by contrast, considers endogenously formed networks and arbitrary neighborhoods (while Acemoglu et al. (2011) consider neighborhoods of the type $K^i \subseteq \{1, 2, \dots, i-1\}$). Other related papers in this literature include Banerjee (1992), Bikhchandani et al. (1992), Bala and Goyal (1998), and Gale and Kariv (2003).

The remainder of this paper is organized as follows. The next section discusses a fairly general way to describe agent-based models and derives a measure for the quality of a hypothesis that is tested in an ABM. Section 3 develops the baseline model and presents the results in the limiting case of an exogenous network structure. Section 4 generalizes the model with exogenous network structure and presents the results for equilibrium network structures and the joint model with uninformative private signals and endogenous network formation. Section (5) concludes.

2 Interpreting Agent-Based Models

Before presenting the model and the main results of this paper, and although this is not a paper about the foundations of agent-based models, it is instructive to have a closer look at the key elements that can be found in any agent-based model. This section gives a rather broad and abstract def-

initiation of an agent-based model and develops a practical criterion to assess the validity of a hypothesis that is tested using an agent-based model.

Agents represent either individuals (e.g. humans) or organizational units (e.g. firms, banks, households, or governments) that can be subsumed in a meaningful manner. An agent a^i is a collection of a set of externally observable actions $x^i \in X^i$, externally non-observable (internal) variables $v^i \in V^i$, exogenously given parameters $p^i \in P^i$, and an information set I^i which contains all information available to the agent. Each agent has a set of neighboring agents whose actions she can observe. The information set thus captures the structure of interactions amongst agents captured in a network structure \mathbf{g} . Internal variables and parameters define the state $s = \{v_i, p_i \forall i\} \in S$ of the world at each point in time t . Agents receive a reward from their actions in a given state of the world which is captured in a reward function R_i . The decision agents take is determined by a policy function π_i which the agent evaluates given an information set I^i and an action x^i to obtain a reward r^i . Agent-based models capture model dynamics in the form of a transition function which describes how the current state of the system changes given the agents' actions.

Definition 1 *An agent-based model (ABM) with $i = 1, \dots, N$ agents a^i is a partially observable Markov decision process consisting of: (i) A space S of states s ; (ii) An action space for each agent i , $\{X_1, \dots, X_i, \dots, X_N\}$; (iii) A transition function: $M : S \times X_1 \times \dots \times X_N \mapsto [0, 1]$ such that:*

$$\int_{S'} M(s, \mathbf{x}, s') ds' = P(s_{t+1} \in S' | s_t = s \text{ and } \mathbf{x}_t = \mathbf{x}),$$

denotes the probability that the vector \mathbf{x}_t of actions of all agents at time t leads to a transition in state space from s_t to $s_{t+1} \in S'$ is some region in S such that $S' \subset S$; (iv) A reward function for each agent i : $R_i : S \times A_1 \times \dots \times A_N \mapsto \mathbb{R}$; (v) A policy function $\pi_i : \tilde{S}_i \times A_1 \times \dots \times A_N \mapsto [0, 1]$, where \tilde{S}_i is the subspace of S observable by agent i , i.e. $\tilde{S}_i = F(s|I_t^i)$, where

F maps the state vector s to the observable state vector of agent i \tilde{s}_i given the information set $I_{i,t}$ available to agent i :

$$\pi_i(\tilde{s}_i, \mathbf{a}) = P(a_{t,i} = a_i | \tilde{s}_t = \tilde{s} \text{ and } a_{tl} = a_l \forall l \neq k).$$

An agent-based model is denoted as $\Gamma = (S, \{A_i\}, M, \{R_i\}, \{\pi_i\}, \{I_i\})$.

The agent's optimization can be computationally intensive and even intractable for Markov decision processes with many agents. Usually, agent-based models are implemented on a computer and solved explicitly in a simulation, which can be defined as:

Definition 2 *An implementation of an ABM is a collection of computer-executable code which contains all elements of an ABM given in Definition 1 and nothing more. A simulation is an ABM together with a set of initial values $\{x_0^i, v_0^i, p_0^i, \mathbf{g}_0\}$ implemented on a computer. The result of a simulation using the parameter set p is the state $\lambda(T; p)$ obtained at the end of the simulation at $t = T$.*

Parameters can be distinguished into model parameters (e.g. a state of nature, denoted θ) and simulation parameters (e.g. the duration T). Using this terminology, a hypothesis that can be tested using an agent-based model is a statement about the result of a simulation:

Definition 3 *A hypothesis of an agent-based model Γ is a result of a simulation $\lambda^H(T; p)$. A parameter-independent hypothesis does not explicitly depend on the parameter set: $\lambda^H(T; p) = \lambda^H(T)$.*

One crucial difference between analytical results and the result of an ABM simulation is that simulation results depend on initial values and parameters. Therefore, a simple single measure is needed to quantify the generality and validity of a hypothesis. One such measure is the *goodness* of a hypothesis, defined as:

Definition 4 *The goodness g of a hypothesis λ^H tested with an ABM Γ is defined as:*

$$g(\lambda^H; \Gamma) = 1 - \frac{\int (\lambda^H(T; p) - \lambda(T; p))^2 dp}{\int dp} \quad (1)$$

where the integral is taken over the entire parameter space (which can be high-dimensional) and each initial value is understood to be a parameter. A goodness value of 1 indicates that the result of the ABM exactly yields the hypothesis for the entire parameter space. Lower goodness indicates that the ABM gives less support to the hypothesis (i.e. the hypothesis is less valid in a smaller parameter space). An agent-based model yields a strong result if it validates a parameter-independent hypothesis with a goodness $g \approx 1$. Two short example illustrate these definitions.

Cournot competition. A simple game that can be formulated as an agent-based model is the Cournot competition game. Two firms indexed by i have to decide their production quantity q_i given a utility function:

$$U_i(q_1, q_2) = p \left(\sum_j q_j \right) q_i - c(q_i),$$

where $p(\cdot)$ and $c(\cdot)$ are the pricing and cost functions respectively. The best response of agent i to the quantity decision of agent $-i$ is given by the reaction function which is obtained from maximising agent i 's utility, $r_i(q_{-i}) = q_i$. A Nash equilibrium is given when $r_1(q_2) = q_1$ and $r_2(q_1) = q_2$ simultaneously. Given certain regularity conditions the Nash equilibrium can be found iteratively, i.e. firms respond optimally to their competitors choice in the previous time step. In this setting we have:

$$q_{it+1} = r_i(q_{-it}).$$

This can be understood as an ABM with a completely observable state space using the following definitions. The vector of state spaces at time t is given by: $s_t = (q_{1t-1}, q_{2t-1})$. The action of agent i is given by $A_{it} = q_{it}$.

The transition function is simply the identity mapping of actions to states; i.e. the state vector at time $t + 1$ is simply the action vector at time t . The reward function is given by u_i . The policy function is given by response function: $\pi_i = r_i(q_{it-1})$.

Foreign exchange trading. Chakrabarti (2000) develops an agent-based model of the foreign exchange market focussing on the endogenous formation of the bid-ask spread offered by dealers in this market. The model reproduces a U-shaped pattern of the bid-ask spread throughout the day - a well documented empirical feature of this market. In Chakrabarti's model dealers estimate the total aggregate demand for a particular currency through stochastic order arrival. Dealers are risk averse utility maximisers and may also request quotes from other dealers to learn about their bid-ask spread. Chakrabarti (2000) deduces a functional form for the optimal bid-ask spread as a function of current inventory levels and the estimated mean and variance of the end of day price of the currency.

The model of Chakrabarti (2000) can be formulated in the language of Definition 1 as follows. The state vector is $s = (\mathbf{Q}, \hat{\mathbf{p}}, \text{var}(\mathbf{p}))$, where \mathbf{Q} is the vector of dealer inventories, $\hat{\mathbf{p}}$ is the vector of the end of day price as estimated by the dealers and $\text{var}(\mathbf{p})$ is the vector of variances of the end of day price as estimated by the dealers. The action vector is $a = (a_k, b_k \forall k)$, where a_k and b_k are the ask and bid prices of agent k respectively. The policy function π_k is given by the expressions Chakrabarti (2000) derived to compute the optimal bid and ask prices. The transition function M specifies the stochastic order arrival to the dealers and the dealers' belief updating mechanism. Thereby it determines how the state vector evolves from t to $t + 1$ given the dealers' bid and ask decisions.

Chakrabarti (2000) tests whether his model reproduces the empirically observed U-shaped pattern of the bid-ask spread - the hypothesis λ^H . He

explores the validity of λ^H by sweeping the parameter space for a total of 729 different parameter combinations and finds that the hypothesis holds in a large part of parameter space.

3 Contagious Herding with Fixed Network Structure

This section develops a baseline model of financial intermediaries that receive a private signal about an underlying state of the world and observe the previous actions of their peers upon which they decide on an optimal investment strategy. The key assumption in this section is that a financial intermediary can only observe an exogenously fixed fraction of his peers which is given in a constant network structure \mathbf{g} . Section 3.1 develops the model, Section 3.2 presents the key result of *contagious herding*, and Section 3.3 discusses a number of extensions to the baseline model. The assumption of a fixed network structure is relaxed in Section 4 where a simple rationale for an endogenous network formation model is introduced.

3.1 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.⁷ There are three model parameters θ , λ , and ρ which are identical for all agents i . By a slight abuse of notation the model parameter θ is sometimes called the state of the world and I assume that it can take two values $\theta \in \{0, 1\}$. I refer to the states $\theta = 1$ as good and $\theta = 0$ as bad. The probability that the world is in state θ is denoted as $\mathbb{P}(\theta)$. 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 is matched by the investment strategy chosen, and nothing otherwise. Agents

⁷I employ a broad notion of financial institutions that encompasses commercial banks, investment banks, money market funds, and hedge funds.

take an action by choosing an investment strategy. Taking an action and switching between actions is costless. The utility of bank i from investing is given as:

$$u^i(x^i, \theta) = \begin{cases} 1 & \text{if } x^i = \theta \\ 0 & \text{else} \end{cases} \quad (2)$$

The state of the world is unknown ex-ante and revealed with probability $\rho \in [0, 1]$ at each point in time t . Once the state is revealed, it can change with probability $\lambda \in [0, 1]$ and if not stated otherwise, $\lambda = \frac{1}{2}$ is assumed which ensures that information about the current state does not reveal information about the future state.⁸ This setup captures a situation where the state of the world, good or bad, is revealed less often (e.g. quarterly) than banks take investment decisions (e.g. daily). In 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

⁸One interpretation is that the system starts from $t = 0$ again once the state of the world is revealed.

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 $s_t^i \in \bar{S}$ where \bar{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 variance (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 $j \in K^i$. On average every $1/\rho$ periods the state of the world is revealed and banks realize their utility. 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 3.2 the state of the world was revealed only at the end of the simulation and only after the system has reached a steady state in which agents do not change their actions any more.⁹

Banks form a private belief at time t based on their privately observed signal

⁹In practice this is ensured by having many more update steps than it takes the system to reach a steady state.

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:

$$I_t^i = \{s_t^i, K_{t-1}^i, x_{t-1}^j \forall j \in K_{t-1}^i\} \quad (3)$$

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 .

Using this nomenclature it is possible to define an equilibrium of the game of social learning with exogenously fixed network structure described in this section:

Definition 5 *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 expected pay-off of bank i given the strategies of all other banks σ^{-i} .*

For every strategy profile σ the expected pay-off of bank i from action $x^i = \sigma^i(I^i)$ is denoted $\mathbb{P}_\sigma(x^i = \theta | I^i)$. Thus for any equilibrium σ , bank i chooses action x^i according to:

$$x^i = \sigma^i(I^i) \in \arg \max_y \mathbb{P}_{(y, \sigma^{-i*})}(y = \theta | I^i) \quad , \quad y \in \{0, 1\} \quad (4)$$

(see Acemoglu et al. (2011)). In their setting, each agent i receives a private signal and observes the actions of a set of neighbors which, by construction, chose their actions before agent i . Agent i then decides on an optimal action. The authors show that there exists a pure strategy perfect Bayesian equi-

librium inductively. Their result carries over to the present setting where each bank receives a private signal and averages over the previous actions of a fixed set of neighbors. The equilibrium is still a Bayesian equilibrium because the averaging over neighbors' actions effectively replaces the entire neighborhood by a single representative agent.¹⁰

The action chosen by bank i carries over from Acemoglu et al. (2011) similarly to the existence of a pure strategy equilibrium and yields:

Proposition 1 *Let σ be an equilibrium of the single bank investment game and let $I_t^i \in \mathcal{I}^i$ be the information set of bank i at time t . Then the strategy decision of bank i , $x_t^i = \sigma^i(I_t^i)$ satisfies*

$$x^i = \begin{cases} 1, & \text{if } \mathbb{P}_\sigma(\theta = 1|s_t^i) + \mathbb{P}_\sigma(\theta = 1|x_{t-1}^j, j \in K_{t-1}^i) > \bar{x}_t \\ 0, & \text{if } \mathbb{P}_\sigma(\theta = 1|s_t^i) + \mathbb{P}_\sigma(\theta = 1|x_{t-1}^j, j \in K_{t-1}^i) < \bar{x}_t \end{cases} \quad (5)$$

and $x^i \in \{0, 1\}$ otherwise.

where the first term on the left-hand side of Equation 5 is the private belief, the second term is the social belief, and where \bar{x}_t is a threshold that has to satisfy two conditions: (i) In the case with no social learning ($K_{t-1}^i = \emptyset$), the threshold should reduce to the simple Bayesian threshold $\bar{x}_B^i = \frac{1}{2}$ 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 threshold 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. A simple function that satisfies both requirements is given by:

$$\bar{x}_t = \frac{1}{2} \left(1 + \frac{k_{t-1}^i}{(n-1)} \right) \quad (6)$$

and yields $\bar{x} \in [\frac{1}{2}, 1]$. Other functional forms are possible, and in particular in the standard model of social learning as employed for example in Gale

¹⁰The action taken by this agent is not a binary action, however. Formally $x_{t-1}^i \in [0, 1]$.

and Kariv (2003), and Acemoglu et al. (2011), each agent observes the action of only *one* neighbor at a time and the threshold does not account for the relative size of the neighborhood. In this case the threshold is simply given by $\bar{x}_t = 1$.

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

$$p^i = \left(1 + \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s_t^i)\right)^{-1} = \left(1 + \frac{f_0(s_t^i)}{f_1(s_t^i)}\right)^{-1} \quad (7)$$

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 by simply averaging over the actions of all neighbors $j \in K_{t-1}^i$:

$$\mathbb{P}_\sigma(\theta = 1|K_t^i, x^j, j \in K_{t-1}^i) = 1/k_{t-1}^i \sum_{j \in K_{t-1}^i} x_{t-1}^j \quad (8)$$

which implies that the bank will choose $x^i = 1$ whenever the sum of private and social belief exceeds the threshold \bar{x} .

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.¹¹

¹¹This 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.

The model in this section can be formulated as an agent-based model using Definition (1). Banks are the agents a^i who can choose one of two actions $x^i \in \{0, 1\}$ yielding the action space X^i . The internal variables are given by the private and social belief of agent i , i.e. $v^i = \{p^i, \mathbb{P}_\sigma(\theta = 1 | K_t^i, x^j, j \in K_{t-1}^i)\}$ and the only exogenously given parameter is the state of the world θ which is identical for all agents i . The internal variables and exogenous parameter together form the state space S . Each agent has an information set I^i given by Equation (3). The reward function R_i is given by the utility (2) of agent i and the policy function π^i by the strategy given in Equation (5). The network structure \mathbf{g} is encompassed in the information set. The transition function specifies how the private and social beliefs of each agent are updated given the actions of all agents and the state of the system, i.e. Equations (7) and (8).¹² When the network structure is exogenously fixed, this ABM is denoted as Γ^{exo} .

3.2 Contagious Herding

Before discussing the full model I build some intuition about the results in useful benchmark cases. Using 6 - 7, an agent i will choose action $x_t^i = 1$, whenever Equation 5 yields:

$$\left(1 + \frac{f_0(s_t^i)}{f_1(s_t^i)}\right)^{-1} + \frac{1}{k_{t-1}^i} \sum_{j \in K_{t-1}^i} x_{t-1}^j > \frac{1}{2} \left(1 + \frac{k_{t-1}^i}{(n-1)}\right) \quad (9)$$

First, consider the (hypothetical) case of a completely uninformative private signal, i.e. $f_0(s_t^i) = f_1(s_t^i)$. The above equation simplifies to:

$$\frac{1}{k_{t-1}^i} \sum_{j \in K_{t-1}^i} x_{t-1}^j > \frac{1}{2} \left(\frac{k_{t-1}^i}{(n-1)}\right) \Leftrightarrow 2(n-1) \sum_{j \in K_{t-1}^i} x^j > (k_{t-1}^i)^2 \quad (10)$$

¹²This ABM is implemented according to Definition 2 using the programming language python. The source code can be found in the supplementary material for this paper at <http://www.co-georg.de/>.

For a maximally connected agent $k_t^i = (n - 1)$ this implies that $2 \sum_j x_{t-1}^j > (n - 1)$ from which it follows that agent i chooses $x_t^i = 1$ whenever more than half of the agents chose $x_{t-1}^j = 1$ in the previous period. A fully connected network with fully uninformative private signals where more than half of the agents selects action $x_t^i = 1$ at any point in time t (e.g. for an exogenous reason or out of pure chance) will thus yield a long-run equilibrium in which $x^i = 1 \forall i$ (and similarly if a majority ever selects $x_t^i = 0$). The long-run equilibrium of the system is thus determined by the initial conditions similarly to the DeGroot (1974) model. Equivalently, a minimally connected agent i with $k_{t-1}^i = 1$ will follow his neighbor in whatever decision the neighbor takes. To see this, consider a star network with $n \geq 3$ nodes where $k^0 = 2$ (the central node is indexed 0). For simplicity assume $x_{t=0}^0 = 1$ and $x_{t=0}^1 = \dots = x_{t=0}^n = 0$. In $t = 1$ the two spokes will follow the central node, while the central node will follow the two spokes and $x_{t=1}^0 = 0$, while $x_{t=1}^1 = \dots = x_{t=1}^n = 1$ and the system always oscillates between those two states, although the vast majority of agents chose a particular action initially. The social belief does no longer carry information about the state of the world, which shows that it's informativeness depends on the network structure.

Second, consider the case of informative private signals. When the network is fully connected, each bank observes the actions of all other banks in the system. Since banks receive only their private signal in the initialization step and since private signals are informative (even if only slightly so), there will be more than half the banks in the system that choose a state-matching action in $t = 0$ and the social signal is state-matching. In the long-run, banks will thus choose a state-matching action whenever the system is fully connected. If, on the other hand, the network is completely empty, banks receive only their private signal at each update step. While the signal is informative, there is no learning from previous signals in this model. Thus, the probability of choosing a state-matching action in period

t depends solely on the private belief in this period. The probability of finding a system in which all banks chose a state-matching action will therefore be higher for more informative signal structures ($\mu_1 - \mu_0 \gg 0$) than for uninformative signal structures ($\mu_1 - \mu_0 \approx 0$).

For interim levels of connectivity the system can enter a *contagious regime* in which all banks choose a state-non-matching action. There exists positive probability that a bank i has a neighborhood K^i which contains more banks that chose a state-non-matching action than banks that chose a state-matching action. The effect can be large enough to offset any private signal bank i might receive. In this case, bank i decides to act on its social belief instead of following its private belief. This contagious process continues in the next period until finally, after a number of update steps, a majority of the financial system follows a misleading social belief instead of their (on average correct) private beliefs. This phenomenon can be seen in Figure (1) in Appendix (A) which shows the average action of all banks at $t = 100$ for varying density of the exogenously fixed random network. The density $\delta(g)$ of the network is defined as

$$\delta(g) = \frac{|g|}{n(n-1)} \quad (11)$$

where $|g|$ is the number of connections in the network given in Equation (14). The density of the network is fixed exogenously and the network is created accordingly.¹³ Each simulation was conducted with $n = 50$ agents and repeated 100 times. Three sets of model parameters for the signal structure are chosen: a highly informative signal structure ($\mu_0 = 0.25, \mu_1 = 0.75, \sigma_{\{0,1\}}^2 = 0.1$), a less informative signal structure ($\mu_0 = 0.4, \mu_1 = 0.6, \sigma_{\{0,1\}}^2 = 0.1$), and a signal structure with very low informativeness ($\mu_0 = 0.49, \mu_1 = 0.51, \sigma_{\{0,1\}}^2 = 0.1$).

¹³The network generation algorithm is very simple: loop over all possible pairs of neighbors in the network and draw a random number. If it is below the exogenously defined network density, add the link to the network and do nothing otherwise. The result of this procedure is a randomly connected network of the requested density.

Figure (1) shows the existence of the contagious regime for low levels of interconnectedness mentioned above. In the contagious regime false beliefs about the state of the world can propagate through the network to all banks in the system. The intuition behind this effect is that the network substitutes a memory and provides banks with a means to receive a signal about previous actions (albeit, of other banks). It can be seen from Figure (1) that the contagious regime is larger for a less informative signal structure (i.e. exists for a larger range of network densities) than for a more informative signal structure.

In order to more rigorously test the initial result discussed above we formulate a null hypothesis which can be tested for robustness using the goodness defined in Definition (4):

Hypothesis 1 *The model of bank behaviour with exogenously given network structure presented in this section and given by Γ^{exo} always leads all agents to coordinate on a state-matching action.*

In accordance with the initial result, a simulation study using the goodness of the hypothesis invalidates the null hypothesis, namely:

Result 1 *The model of bank behaviour with exogenously given network structure presented in this section and given by the ABM Γ^{exo} exhibits a contagious regime for interim levels of interconnectedness in which all agents coordinate on a state-non-matching action.*

Goodness as a function of network density for Γ^{exo} is shown in Figure 2 for $\theta = 0$ and varying degrees of network density. The parameters are the signal informativeness (i.e. the difference between the mean of \mathbb{F}_1 and \mathbb{F}_0), as well as the signal variance $\sigma_{\{0,1\}}^2$. Signal informativeness varies between 0.02 and 0.9 in 0.1 steps, and signal variance between 0.02 and 0.25 in 0.05 steps. Figure 2 shows that result 1 holds true and quantifies the interim range of network density to be smaller than 0.2. Each point is the result

of 1000 simulations with 50 sweeps and $N = 15$ agents. The results hold qualitatively when the number of agents is varied.

3.3 Extensions of the Baseline Model

After establishing the main result for the baseline model, this section presents two interesting extensions that explore various aspects of agent heterogeneity. First, heterogeneity in terms of *interconnectedness* refers to the possibility that some agents are significantly better connected and therefore observe significantly more actions of neighbors. Second, heterogeneity in terms of *informedness* refers to the fact that some agents might receive significantly better private signals.

Different Network Topologies. Interbank networks can exhibit a range of possible network structures.¹⁴ Three types of networks can be distinguished: (i) Random networks where the probability that two nodes are connected is independent of other characteristics of the node; (ii) Barabasi-Albert (or scale-free) networks where the probability that a new node entering the system is connected to an existing node is proportional to the existing node's degree; (iii) Watts-Strogatz (or small-world) networks which are essentially regular networks with a number of “shortcuts” between remote parts of the network. Three parameters are useful to differentiate between the different types of networks: network density (number of links divided by the number of possible links), shortest average path length (number of links between two random nodes), and clustering (the probability that a node's counterparties are counterparties of each other. In random networks clustering and shortest average path length is proportional to the density of the network. In Barabasi-Albert networks some nodes have substantially higher clustering than the rest of the system, even for low values of network density (and thus large shortest average path length). And Watts-Strogatz

¹⁴See Georg (2013) for an overview of empirical works on the interbank network structure and a comparison of their susceptibility to financial contagion.

networks feature low shortest-average path lengths even for relatively high values of clustering.

Figure 3 shows the results for the goodness of the hypothesis that banks coordinate on a state-matching action for $\theta = 0$ for various network topologies. Random networks are created using $\rho = \{0.1, 0.2, 0.3, 0.4, 0.5\}$, Barabasi-Albert networks with each new node being connected to $k = \{2, 3, 4, 5, 6, 7, 8, 9, 10\}$ others (Barabasi-Albert networks are constructed, starting with k initial nodes and nodes being added until the network has N nodes), and Watts-Strogatz networks (which are constructed from a regular graph in which each node has m neighbors and the probability that a given link is attached to a random node is p) for $m = \{2, 4, 8\}$ and $p = \{0.05, 0.1, 0.2, 0.3, 0.5\}$.

The goodness of the hypothesis that banks choose a state-matching action is lower for larger values of the shortest average path length for all types of networks. While Barabasi-Albert networks are characterized by a small number of very highly connected banks and a large number of less interconnected banks, this does not improve the goodness of the hypothesis if the average path-length is too large. Note that the drop in goodness occurs at the same value of average shortest path-length suggesting that the effect is independent of the exact network structure. Similarly, lower values of the clustering coefficient are associated with low values of goodness, however, not identically for all network topologies.

Agent Heterogeneity. A simple way to capture agent heterogeneity is to introduce two types of agents, informed and uninformed. Informed agents have a much higher informativeness of their private signal (i.e. $\mu_1 - \mu_0 \gg 0$) than uninformed agents (i.e. $\mu_1 - \mu_0 \approx 0$). Heterogeneity is introduced by varying the probability that a fraction p_{inf} is informed and $(1 - p_{\text{inf}})$ is uninformed. Figure 4 shows the result of varying agent heterogeneity on the goodness of the hypothesis that the model exhibits a contagious regime.

The probability p_{inf} is increased from 0 to 1 in 0.1 steps and the goodness is computed for each value of p_{inf} separately. Being informed corresponds to mean signals of $\mu_0^{\text{inf}} = 0.25$, $\mu_1^{\text{inf}} = 0.75$, and $\left(\sigma_{\{0,1\}}^2\right)^{\text{inf}} = 0.1$ while being uninformed corresponds to mean signals of $\mu_0^{\text{!inf}} = 0.49$, $\mu_1^{\text{!inf}} = 0.51$, and $\left(\sigma_{\{0,1\}}^2\right)^{\text{!inf}} = 0.1$. I consider two cases, one where the network structure would exhibit a contagious regime, i.e. where the density is 0.1 and one where there is no contagious regime any more, i.e. where the density is 0.5.

In both cases the goodness does not change much as a function of the probability of an agent being informed. For the non-contagious network density, almost no variation in goodness is observed, while the simulations with low density show little variation in the goodness of the hypothesis. The special case of only informed agents is also shown in Figure 1, confirming that the range of low network density ($\rho \approx 0.1$) indeed exhibits the contagious regime, even if some agents are more informed than others. The intuition is simple, even well-informed agents will eventually act on their social signal instead of their private signal if sufficiently many neighbors act uniformly.

4 Endogenous Network Formation

The assumption of an exogenously fixed network structure is rather restrictive for financial networks as studies of money market network structures show (see, for example, Arciero et al. (2013), and Gabrieli and Georg (2013) for an analysis of the European interbank market). This section therefore relaxes this assumption and develops a simple model of endogenous network formation in which banks mutually decide on which links to form.

A few modifications to the model with exogenous network structure are in order, in particular the time-line is slightly modified. In $t = 0$ there is no endogenously formed link and each bank decides on its action in autarky. Banks receive a signal about the state of the world and form a private be-

lie upon which they decide about their investment strategy x^i . In $t = 1$ and all subsequent periods banks receive their private signal and also observe the actions by all neighbors in the previous period. Based upon this information, banks decide on their action. Once banks have chosen their individual actions they agree on a new network of mutual lines of credit. An equilibrium outcome for this game will be obtained by using the notion of pairwise stability introduced by Jackson and Wollinsky (1996). In order to utilize pairwise stability, a bank i 's benefits and risks of being connected to bank j have to be determined.

In this simple extension, I assume three motifs for banks to engage in inter-bank lending. First, banks have a benefit from forming a connection since they receive additional information about the state of the world. However, forming and maintaining a link is costly due to increased business operations (e.g. infrastructure cost, cost of operating a trading and risk management department). The net benefit of establishing a link between banks i and j is given as

$$\alpha g^{ij} \tag{12}$$

where $\alpha \in \mathbb{R}$ is positive unless specified otherwise. The benefit of learning a neighbor's action is thus assumed to outweigh the costs of establishing and maintaining a link. As long as agents coordinate on a state-matching action, learning a neighbor's action is additional and valuable information, increasing the probability of selecting a state-matching action. In this case, it is possible to compute the value from learning in a closed form. However, agents do not know when their neighbors have selected a state non-matching action. Thus, a closed form solution for α for all cases is infeasible. Absent any other motif, a positive α for all banks implies that the resulting network structure will be perfectly connected, an observation at odds with observed interbank network structures. Therefore, either α is stochastic (positive for some, but not for all banks), or other motifs for interbank network formation must be present. In this paper I follow the second approach.

Second, when the state of the world is revealed a bank i that did not choose the correct strategy ($x^i \neq \theta$) will incur a liquidity shortfall. This can be motivated by the following argument. Assume the state of the world is a bust and that bank i has chosen a strategy that yields a portfolio which performs good in a boom, but badly in a bust. Once the state of the world is revealed there is a probability of $\lambda = 1/2$ that the state of the world changes and remains in the new state for (on average) $1/p$ periods. In such a situation a bank that chose a non-state-matching action will suffer liquidity outflows as investors will try to put their money in banks with better adjusted portfolio.¹⁵ A mutual line of credit between banks i and j implies that bank i can draw upon liquidity from bank j and avoid costly fire-sales. Thus, there exists a positive probability that bank j chose the state-matching action and i can draw upon the mutual line of credit and avoid a fire-sale whenever $x^i \neq \theta$. When bank i has a private belief of $p^i = 1/2$ it is, without additional information, completely uncertain about the underlying state of the world and coinsurance will be most valuable. If bank i is certain about the state θ of the world, it will not value coinsurance at all. This is captured by introducing a value-function $q^i(p^i)$ defined as:

$$q^i(p^i) = \begin{cases} 2p^i & \text{for } p^i \leq \frac{1}{2} \\ 2(1 - p^i) & \text{for } p^i \geq \frac{1}{2} \end{cases} \quad (13)$$

which can be used to define the expected utility from coinsurance:

$$\beta_1 q^i(p^i) g^{ij} \quad (14)$$

where $\beta_1 \in \mathbb{R} > 0$.¹⁶

¹⁵The underlying assumption is that portfolios are more liquid when matching the state of the world. Banks that chose the right strategy can then adjust their portfolio more quickly to a (possible) new state of the world than banks that are stuck with an illiquid, state-non-matching portfolio.

¹⁶One could argue that the value of additional links is declining in the total number of links. Since banks average over all their neighboring links with equal weights, however, each link carries equal value.

While the upside from mutual lines of credit is liquidity coinsurance, the downside is counterparty risk which arises whenever a bank i chose a state-matching strategy and is connected to a bank j which chose a non-state-matching strategy.¹⁷ Counterparty risk leads to losses due to contagious defaults. The intuition is similar to that of Equation (4): whenever bank i assumes to be right about the state of the world there is a risk that bank j is not right in which case i will incur a loss due to contagion. This is captured by assuming that the expected loss from counterparty risk is given as:

$$- \beta_2 (1 - q^i(p^i)) g^{ij} \quad (15)$$

where $\beta_2 \in \mathbb{R} > 0$.¹⁸ For simplicity I assume $\beta_1 = \beta_2$ in which case the benefit from coinsurance and the expected loss from counterparty risk sum up to:

$$\beta(2q^i(p^i) - 1)g^{ij} \quad (16)$$

which has a natural interpretation. When bank i is certain about the state of the world, it will fear contagion more than the benefit from coinsurance and Equation (4) is negative. While, on the other hand, the benefit from coinsurance will outweigh the potential loss from counterparty risk whenever bank i is uncertain about the state of the world.

Finally, to capture amplification effects from financial fragility I assume that banks suffer an expected loss of:

$$- \gamma |g| q^i(p^i) g^{ij} \quad (17)$$

¹⁷In Castiglionesi and Navarro (2011) both liquidity coinsurance and counterparty risk is considered in a microfounded model of banking behaviour.

¹⁸Counterparty risk in interbank markets can lead banks to cut lending and eventually even to money market freezes. For theoretical contributions see Rochet and Tirole (1996), and Heider et al. (2009), and Acharya and Bisin (2013). Empirically the role of counterparty risk in market freezes has been analyzed for example by Afonso et al. (2011).

where $\gamma \in \mathbb{R} > 0$.¹⁹ The term will increase with the total number of connections in the financial system, given as

$$|g| = \sum_{i,j} g^{ij} \quad (18)$$

This amplification effect captures a situation where a number liquidity constraint banks are forced to sell assets in a fire-sale. In essence, it captures a non-linear downward-sloping demand for assets similar to Cifuentes et al. (2005). I will analyze the equilibrium network structures that are obtained from these three motifs below.

Bank i 's utility from forming a connection with bank j is thus given as:

$$u^i(g^{ij} = 1) = \alpha + \beta (2q^i(p^i) - 1) - \gamma |g| q^i(p^i) \quad (19)$$

The utility bank i receives from being interconnected is

$$u^i(g) = \sum_{j \in K_i} u^i(g^{ij} = 1) \quad (20)$$

and the utility of the whole network $u(g)$ is $u(g) = \sum_i u^i(g)$.

An update step consists of agents choosing an optimal strategy based on their private and social beliefs, and a network formation process. Following Jackson and Wollinsky (1996) an equilibrium of the network formation process can be characterized using the notion of pairwise stability.

Definition 6 *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 $u^j(g) \geq u^j(g - l^{ij})$*

¹⁹Korinek (2011) shows how such amplification effects can arise in a framework of financial fragility.

- (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.

The notion of a pairwise stable equilibrium describes a situation where two banks that both obtain positive utility from establishing a mutual link will do so, and all others will not. This cooperative notion of equilibrium is only one possibility, however. Bala and Goyal (2000) develop an equilibrium concept based on non-cooperative network formation where the cost of a link is borne by one of the banks only. Recently, Acemoglu et al. (2013) and Gofman (2013) extend the approach by Eisenberg and Noe (2001) to describe the formation of networks. Analysing the impact of such alternative network formation concepts are beyond the scope of the present paper, however.

4.1 The Structure of Endogenous Networks

While Section 3.2 considered the special case of a fixed network structure and not highly informative signal, this section does the converse. It describes a situation in which each bank receives a highly informative private signal about the underlying state of the world and forms links endogenously. This is achieved by using $\mu_0 = 0.25$, $\mu_1 = 0.75$, and $\sigma_{\{0,1\}}^2 = 0.05$. I furthermore assume that there exist $i = 1, \dots, n$ ex-ante identical banks, i.e. there is no agent heterogeneity. The network structure is formed endogenously depending on the three parameters α , β , and γ in Equation (15). Each parameter captures a different motive to form or sever an interbank link: $\alpha \in [0, 1]$ describes the utility that is obtained through enhanced learning; $\beta \in [0, 1]$ is the coinsurance-counterparty risk trade-off term capturing the fact that banks who do not have a very informative private signal about the state of the world prefer to form connections to reap the benefits of coin-

surance (while banks with a very informative private signal prefer to not form a link to avoid counterparty risk); and γ accounting for amplification effects in the form of, for example, fire-sales.

Two trivial results can be readily observed. First, if $\beta = \gamma = 0.0$, any positive value of α leads to a complete network, while $\alpha = 0$ yields an empty network. And second, for $\alpha = \beta = 0.0$, any value of γ yields an empty network (recall that a link is only formed if both agents obtain positive utility from the formation of the link).

For $\alpha = \gamma = 0.0$ and $\beta > 0.0$ bank i 's utility from establishing a link is given as $\beta (2q^i(p^i) - 1)$ and thus positive whenever $q^i(p^i) > 1/2$. This happens if $p^i \in (1/4, 3/4)$.²⁰ For very uninformative signal structures ($\mu_1 - \mu_0 \approx 0$) agents almost never receive a sufficiently strong private signal so they would find it optimal to not form a link. The resulting network structure is that of a complete network. For very informative signal structures ($\mu_1 - \mu_0 \approx 1$) agents are almost always certain about the state of the world from just using their private signal and will find it thus almost never beneficial to establish a link. The resulting network is empty. The network density will be between those two extremes for interim ranges of informativeness. This intuition is confirmed by Figure (5) which was obtained by measuring the equilibrium network density (i.e. the network structure in $t = 20$) for different levels of signal informativeness ($\mu_1 - \mu_0 \in [0.02, 1.0]$). Each simulation consists of $n = 20$ agents and was repeated 500 times. The network density shown in Figure (5) is the average network density over all 500 simulations.

Another interesting limiting case is obtained whenever there is a positive utility from ($\alpha > 0$) learning but negative utility from amplification effects ($\gamma > 0$). Banks receive a constant utility from learning, while the disutility from amplification effects grows with the number of links in the network.

²⁰This follows immediately from the definition of q^i in Equation (9).

Once the number of links is too large, no further links are added. The resulting network structure is that of a star which is characterized by (i) small average shortest path length $l \simeq 2$; (ii) density $\delta = 1/n$; and (iii) a clustering coefficient of zero. The average shortest path length is defined as:

$$l(i, j) = \sum_{i, j} \frac{d(i, j)}{n(n-1)} \quad (21)$$

where the sum runs over all nodes i, j and where $d(i, j)$ is the length of the shortest path from i to j . The clustering coefficient $c(g)$ is defined via the local clustering coefficient c^i of bank i :

$$c^i = \frac{|\{l^{jk}, j, k \in K^i, l^{jk} \in L\}|}{k^i(k^i - 1)} \quad (22)$$

as:

$$c = \frac{1}{n} \sum_i c^i \quad (23)$$

Intuitively, the local clustering coefficient of a bank i is the probability that two neighbors of i are connected.

Figures (6) and (7) show the three network measures as a function of the amplification parameter $\gamma \in [0, 1]$ with fixed learning parameter $\alpha = 0.01$. Each point is the result of 500 simulations with $n = 20$ agents. For a perfect star with n nodes, $l \simeq 2$, $c = 0.0$, and $\delta = 1/n$. For the uninformative signal structure $\mu_0 = 0.4, \mu_1 = 0.6$ shown in Figure (6) a star-like network is obtained for $\gamma > 0.3$. For the informative signal structure $\mu_0 = 0.25, \mu_1 = 0.75$ shown in Figure (7) a star-like network is much less pronounced and the resulting network will be a mixture of a star and a random network.

From the discussion above it can be seen that the equilibrium network outcomes are a superposition of the limiting cases. The trade-off between coinsurance and counterparty risk (as a result from signal informativeness) controls the overall network density. The trade-off between learning and

amplification controls how star-like the equilibrium network will be. This can be summarized as follows:

Result 2 *The density of the pairwise stable equilibrium network that is obtained in the pure coinsurance-counterparty risk case ($\alpha = \gamma = 0, \beta > 0$) decreases with the informativeness of the signal structure. The pairwise stable equilibrium network that is obtained with positive utility from learning in the presence of amplification effects ($\alpha > 0, \gamma > 0, \beta = 0$) is star-like.*

4.2 Learning and Endogenous Network Formation

Section (3.2) shows the existence of a contagious regime for varying signal informativeness and low network density. In this regime there is a positive probability that banks synchronize on a state-non-matching action. The contagious regime is larger when the signal structure is less informative. At the same time, however, Section (4.1) shows that the density of an endogenously formed network is large for low signal informativeness since banks will put more emphasis on the liquidity coinsurance character of mutual lines of credit. This raises the question which of the two effects will prevail when they are both present in a model with learning and endogenous network formation.

To address this question, Figure (8) shows the total utility of agents in the system as a function of signal informativeness. Total utility U is defined as the sum of two terms:

$$U = \sum_i [u^i(x^i, \theta) + u^i(g)] \quad (24)$$

where the first term is the individual utility bank i receives from choosing a state-matching action and the second term is the utility bank i receives from being interconnected $u^i(g) = \sum_{j \in K^i} u^i(g^{ij} = 1)$. Each point in the figure is the average utility from 500 simulations with $n = 20$ agents and has been measured at the end of each simulation in $t = 20$. The individual

utility will always be in the range $u^i(x^i, \theta) \in [0.0, n]$ and the relation between individual utility and utility from interconnectedness is controlled by the parameter β . Larger values of β imply relatively larger values of utility from interconnectedness.

Figure (8) shows the trade-off between social learning in the contagious regime and liquidity coinsurance. For very uninformative signals ($\mu_1 - \mu_0 \simeq 0$) synchronization on state-non-matching actions in the contagious regime exists for larger ranges of network densities and the only utility obtained stems from being interconnected. As signal informativeness increases, the size of the contagious regime gets smaller but at the same time the network density is reduced which increases the chance of being in the contagious regime. The effect from the reduction of the size of the contagious regime, however, is stronger than the effect from reduced network density, as can be seen from Figure (8) for varying strengths of β . The total utility, however, is reduced proportional to the reduction in network density.

This leads to:

Result 3 *In the full model with social learning and endogenous network formation an increasing informativeness of the signal structure the reduction in the size of the contagious regime outweighs the reduction in network density. Unless the signal informativeness is extremely low banks synchronize their strategies on state-matching actions.*

5 Conclusion

This paper develops a model of ex-ante 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. Three aspects determine the equilibrium network structure:

(i) the benefit from learning the signal of counterparties; (ii) a trade-off between coinsurance and counterparty risk: banks with more informative private signal have less incentives to form a link since the coinsurance motif is dominated by counterparty risk; and (iii) the threat of amplification effects when a bank chooses a state-non-matching action.

Three results are obtained. First, I show the existence of a contagious regime in which banks synchronize their investment strategy on a state-non-matching action. This regime is larger for less informative signals and exists in incomplete networks, irrespective of the network type and also for heterogenously informed agents. When a bank is connected to a set of counterparties that on average selected a state-non-matching action the social signal can outweigh the private signal and the synchronization on a state-non-matching action becomes contagious. Second, I characterize the equilibrium interbank network structures obtained. For more informative private signals the network structure becomes sparser since banks fear counterparty risk. For stronger amplification effects star-like networks emerge and the equilibrium interbank network structures obtained in the full model resemble real-world interbank networks. Third, I show that for low signal informativeness the contagious regime still exists, i.e. that the effect from contagious synchronization outweighs the effect from increased network connectivity.

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. This benefit will depend on whether or not a neighboring bank chose a state matching on state non-matching action in the previous period. 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., Dahleh, M.A., Lobel, I., Ozdaglar, A., 2011. Bayesian learning in social networks. *Review of Economic Studies* 78, 1201–1236.
- Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2013. Systemic Risk and Stability in Financial Networks. mimeo. MIT.
- Acharya, V., Bisin, A., 2013. Counterparty risk externality: Centralized versus over-the-counter markets. *Journal of Economic Theory* 149.
- Acharya, V.V., Gale, D., Yorulmazer, T., 2011. Rollover risk and market freezes. *The Journal of Finance* 66.
- Acharya, V.V., Merrouche, O., 2013. Precautionary hoarding of liquidity and the interbank markets: Evidence from the sub-prime crisis. *Review of Finance* 17.
- Acharya, V.V., Skeie, D., 2011. A model of liquidity hoarding and term premia in interbank markets. *Journal of Monetary Economics* 58.
- Afonso, G., Kovner, A., Schoar, A., 2011. Stressed, not frozen: The federal funds market in the financial crisis. *The Journal of Finance* 66.
- Alevy, J.E., Haigh, M.S., List, J., 2007. Information cascades: Evidence from an experiment with financial market professionals. *The Journal of Finance* 67.
- Allen, F., Babus, A., 2009. Networks in finance, in: Kleindorfer, P.R., Wind, Y.R., Gunther, R.E. (Eds.), *Network Challenge, The: Strategy, Profit, and Risk in an Interlinked World*. Pearson Prentice Hall, pp. 367–382.
- Allen, F., Gale, D., 2000. Financial contagion. *Journal of Political Economy* 108, 1–33.
- Arciero, L., Heijmans, R., Huever, R., Massarenti, M., Picillo, C., Vacirca, F., 2013. How to measure the unsecured money market? *The Eurosys-*

- tem's implementation and validation using TARGET2 data. Working Paper 369. De Nederlandsche Bank.
- Babus, A., 2011. Endogenous Intermediation in Over-the-Counter Markets. Working Paper. Imperial College London.
- Babus, A., Kondor, P., 2013. Trading and Information Diffusion in Over-the-Counter Markets. Working Paper. Imperial College London.
- Bala, V., Goyal, S., 1998. Learning from neighbors. *The Review of Economic Studies* 65.
- Bala, V., Goyal, S., 2000. A noncooperative model of network formation. *Econometrica* 68.
- 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.
- 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.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100.
- Bluhm, M., Faia, E., Krahnen, J.P., 2013. Endogenous Banks' Networks, Cascades, and Systemic Risk. SAFE Working Paper. Goethe University Frankfurt.

- Caballero, R., 2012. A Macro Tragedy Waiting to Happen (in a VaR sense). Keynote Speech, TCMB Conference 2012. TURKIYE CUMHURIYET MERKEZ BANKASI.
- Castiglionesi, F., Navarro, N., 2011. Fragile Financial Networks. mimeo. Tilburg University.
- Castiglionesi, F., Wagner, W., 2013. On the efficiency of bilateral interbank insurance. *Journal of Financial Intermediation* 22, 177–200.
- Chakrabarti, R., 2000. Just another day in the inter-bank foreign exchange market. *Journal of Financial Economics* 56.
- 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.
- Chiang, T.C., Zheng, D., 2010. An empirical analysis of herd behavior in global stock markets. *Journal of Banking and Finance* 34.
- Cifuentes, R., Shin, H.S., Ferruci, G., 2005. Liquidity risk and contagion. *Journal of the European Economic Association* 3, 556–566.
- Cohen-Cole, E., Patacchini, E., Zenou, Y., 2012. Systemic Risk and Network Formation in the Interbank Market. mimeo. University of Maryland.
- 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.
- Eisenberg, L., Noe, T.H., 2001. Systemic risk in financial systems. *Management Science* 47, 236–249.

- 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.
- Gabrieli, S., Georg, C.P., 2013. A Network View on Interbank Liquidity. mimeo. Deutsche Bundesbank.
- Gai, P., Haldane, A., Kapadia, S., 2011. Complexity, concentration and contagion. *Journal of Monetary Economics* 58.
- Gai, P., Kapadia, S., 2010. Contagion in financial networks. *Proceedings of the Royal Society A* 466, 2401–2423.
- Gale, D., Kariv, S., 2003. Bayesian learning in social networks. *Games and Economic Behaviour* 45.
- Gale, D., Yorulmazer, T., 2013. Liquidity hoarding. *Theoretical Economics* 8.
- Georg, C.P., 2013. The effect of the interbank network structure on contagion and common shocks. *Journal of Banking and Finance* 37.
- Gofman, M., 2013. A Network-Based Analysis of Over-the-Counter Markets. mimeo. University of Wisconsin Madison.
- Haldane, A.G., May, R.M., 2011. Systemic risk in banking ecosystems. *Nature* 469, 351–355.
- Heider, F., Hoerova, M., Holthausen, C., 2009. Liquidity hoarding and interbank market spreads: The role of counterparty risk. Working Paper 1126. European Central Bank.
- Iori, G., Jafarey, S., Padilla, F., 2006. Systemic risk on the interbank market. *Journal of Economic Behavior & Organization* 61, 525–542.
- Jackson, M.O., Wollinsky, A., 1996. A strategic model of social and economic networks. *Journal of Economic Theory* 71.

- Korinek, A., 2011. Systemic Risk-Taking, Amplification Effects, Externalities, and Regulatory Responses. Working Paper 1345. European Central Bank.
- 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.
- Rochet, J.C., Tirole, J., 1996. Interbank lending and systemic risk. *Journal of Money, Credit and Banking* 28, 733–762.

A Figures




Figure 1: Average actions of agents in $t = 100$ for $\theta = 0$, varying network densities, and different signal structures: (i) high informativeness, $\mu_0 = 0.25, \mu_1 = 0.75, \sigma_{\{0,1\}}^2 = 0.1$; and (ii) low informativeness, $\mu_0 = 0.4, \mu_1 = 0.6, \sigma_{\{0,1\}}^2 = 0.1$; (iii) very low informativeness, $\mu_0 = 0.49, \mu_1 = 0.51, \sigma_{\{0,1\}}^2 = 0.1$. Each point is the average of 100 simulations with $n = 50$ agents.




Figure 2: Goodness of the hypothesis that agents coordinate on a state-matching action for varying network densities for $\theta = 0$. Signal distance ($\mu_1 - \mu_0$) has been varied between 0.02 and 0.9 in 0.1 steps and signal variance between 0.02 and 0.25 in 0.05 steps. For each point, 1000 simulations with 50 sweeps and 15 agents were performed. Goodness computed via the mean of actions is given by the solid line, while dashed lines indicate the goodness result for mean \pm one standard deviation.




Figure 3: Goodness of the hypothesis that agents coordinate on a state-matching action for different network topologies and $\theta =$. Signal distance ($\mu_1 - \mu_0$) has been varied between 0.02 and 0.9 in 0.1 steps and signal variance between 0.02 and 0.25 in 0.05 steps. For each point, 1000 simulations with 50 sweeps and 15 agents were performed. Goodness computed via the mean of actions is given by the solid line, while dashed lines indicate the goodness result for mean \pm one standard deviation. Top row: random networks; Center row: Barabasi-Albert networks; Bottom row: Watts-Strogatz networks; Left column: network density; Middle column: average shortest path length; Right column: clustering.




Figure 4: Goodness of the hypothesis that agents coordinate on a state-matching action for different network topologies and $\theta = 0$. Signal distance ($\mu_1 - \mu_0$) has been varied between 0.02 and 0.9 in 0.1 steps and signal variance between 0.02 and 0.25 in 0.05 steps. For each point, 1000 simulations with 50 sweeps and 15 agents were performed. Goodness computed via the mean of actions is given by the solid line, while dashed lines indicate the goodness result for mean \pm one standard deviation. Goodness is shown for varying probabilities of agent informedness.




Figure 5: Average equilibrium network density as a function of signal informativeness ($\mu_1 - \mu_0$) for $\alpha = \gamma = 0.0, \beta = 0.1$. Each point is the average of 500 simulations with $n = 20$ agents and $\sigma_{\{0,1\}}^2 = 0.1$.




Figure 6: Network density, average clustering coefficient, and average shortest path length for varying amplification parameter $\gamma \in [0.0, 1.0]$ and fixed learning parameter $\alpha = 0.01$. Each point is the average of 500 simulations with $n = 20$ agents and $\mu_0 = 0.4, \mu_1 = 0.6, \sigma_{\{0,1\}}^2 = 0.1$.




Figure 7: Network density, average clustering coefficient, and average shortest path length for varying amplification parameter $\gamma \in [0.0, 1.0]$ and fixed learning parameter $\alpha = 0.01$. Each point is the average of 500 simulations with $n = 20$ agents and $\mu_0 = 0.25, \mu_1 = 0.75, \sigma_{\{0,1\}}^2 = 0.1$.




Figure 8: Total utility (individual + network) as a function of signal informativeness ($\mu_1 - \mu_0$) for $\alpha = \gamma = 0.0, \beta = \{0.02, 0.1, 0.4\}$. Each point is the average of 500 simulations with $n = 20$ agents and $\sigma_{\{0,1\}}^2 = 0.1$.