Reputational Learning and Network Dynamics ☆

Simpson Zhang[1], Mihaela van der Schaar[2]

Abstract

This paper is the first to study reputational learning in networks, taking into account that agents differ from each other, begin with incomplete information, and must learn through past experiences which connections to form and which to break. In our model, agents providing high benefits develop high reputations and remain in the network, while agents providing low benefits will drop in reputation and become ostracized. We characterize the set of stable networks that develop through this learning process. We also use our model to compute the ex ante social welfare of different networks and characterize socially optimal network structures. When the learning speed is either very fast or very slow, we show that complete networks are optimal. When learning is intermediate, other structures such as core-periphery structures can become optimal. We also show that subsidizing links can increase overall welfare by fostering more sustained interactions.

Keywords: Network Dynamics, Network Design, Reputational Learning, Social Welfare

☆The authors are grateful to Jie Xu and Peng Yuan Lai for excellent research assistance. The authors would also like to thank William Zame and Moritz Meyer-ter-Vehn for valuable comments regarding this paper. This paper also benefited from discussions with Matt Jackson and seminar participants at the 2015 NEGT Conference and the 2015 SWET conference. The authors gratefully acknowledge financial support from ONR grant N00014-17-1-2125.

Email addresses: `simpson.zhang@ofr.treasury.gov` (Simpson Zhang), `mihaela.vanderschaar@eng.ox.ac.uk` (Mihaela van der Schaar)

[1]Simpson Zhang is with the Office of Financial Research, U.S. Department of the Treasury. This paper was started while Simpson Zhang was at UCLA; the views and opinions expressed are those of the authors and do not necessarily represent official positions or policies of the Office of Financial Research or the U.S. Department of the Treasury.

[2]M. van der Schaar is with University of California, Los Angeles, CA, USA and the Oxford-Man Institute, University of Oxford, Oxford, UK.

1. Introduction

Networks are pervasive in all areas of society, ranging from financial networks to organizational networks to social networks. And a defining feature of many real world networks is that agents do not fully know the characteristics of others initially and must learn about them over time. For instance a bank learns about the credit-worthiness of a new borrower, a worker in a firm learns about the ability of a coworker, and a buyer learns about the product quality of a supplier. Such learning can strongly affect the resulting shape of the network. As agents receive new information, they revise their beliefs about other agents, update their linking decisions, and cause the network to evolve as a result. For a proper analysis of such networks, it is crucial to understand the exact mechanism by which learning impacts network dynamics.

The impact of agent learning on network evolution has not been well studied in the existing literature. A large network science literature analyzes the effect of learning on fixed networks that have already formed (see [1]). A smaller microeconomics literature[3] studies the formation of networks - but makes very strong assumptions (e.g., homogeneous agents/entities, complete information about other agents). Neither the network science literature nor the microeconomics literature has so far taken into account that agents behave strategically in deciding what links to form/maintain/break *and* that they also begin with incomplete information about others, so they must learn about others through their interactions. As a result, neither network science nor microeconomics provides a complete framework for understanding and guiding the formation (and evolution) of real networks and predicting the consequences of network formation.

As a motivating example of the impact of learning and reputational forces, consider a group of financial institutions that are linked together in a financial network[4]. These financial institutions provide benefits to each other by engaging

[3] See the overview in [2] for instance.

[4] Our model can also be applied to a wide range of other networks, such as organizational

in mutually beneficial trading opportunities, such as providing each other with liquidity or investing in joint ventures[5]. High quality institutions are likely to develop a high realized quality of assets from these joint interactions, while low quality institutions will likely develop a low realized quality of assets. Each institution only continues to link with another institution (over time) if the counterparty is believed to be of sufficiently high quality. As time progresses, the institutions *observe* the actions of their counterparties, *update* their beliefs about the quality of each counterparty, and *change* their linking decisions as a result. In this way, learning by the financial institutions causes the network topology to evolve over time. The network topology also impacts the rate of learning, as an institution can also learn by monitoring interactions of a counterparty with its other counterparties. Institutions with more connections interact with more counterparties, and will therefore reveal more information about themselves to their neighbors over time. As a result, while having more connections opens an institution up to more beneficial opportunities, it also carries the risk of causing the institution to be shut out of the financial network more quickly if it starts losing suffering losses, as in the case of Lehman Brothers due to its exposures to the subprime mortgage market during the 2008 financial crisis.

Our model takes into account the features of the previous example: agents behave strategically, begin with incomplete information about each other, and must learn through continued interactions which connections to form and maintain and which to break. We consider a continuous time model with a group of agents who are linked according to a network and who send noisy flow benefits to their neighbors. The benefits that agents provide could be interpreted for example as the benefits that financial institutions derive from providing liquidity to each other or from diversifying risk with each other's specialized assets. Each agent is distinguished by a fixed quality level which determines the av-

networks, social networks, or expertise networks. We discuss some implications for these settings as well throughout the paper.

[5] As in the model of [3].

erage value of the flow benefits it produces. Agents observe all the benefits that their neighbors produce, and they update their beliefs about a neighbor's quality via Bayes rule. Neighbors with more connections will reveal more information about themselves over time. Agents will maintain links with neighbors that provide high benefits, but will cut off links with neighbors that provide low benefits. The network evolves as agents learn about each other and update their beliefs. Since the number of links an agent has affects the rate of learning about that agent, the rate of learning about an agent changes as the network changes, leading to a co-evolution of network topology and information production.

Our model is highly tractable and allows us to completely characterize network dynamics and give explicit probabilities that the network evolves into various configurations. In addition, we are able to characterize the entire set of possible stable networks and analytically compute the probability that any single stable network emerges. This allows for predictions regarding which stable networks are likely to emerge given an initial network.

We also study the implications that learning has on the social welfare and efficiency of a network. Our results show that learning has a beneficial aspect: agents whose true quality is low are likely to produce low signals and will eventually become ostracized from the network. Learning also has a harmful aspect: even agents whose true quality is high may produce an unlucky string of bad signals and so be forced out of the network. Moreover, even having low quality agents leave the network can reduce overall social welfare. A marginally low quality agent may harm its neighbors slightly, but it also receives a large benefit if its neighbors' qualities are very high. Therefore, if the low quality agent leaves the network the overall social welfare would actually decrease. The issue here is that agents only care about the benefit their neighbors are providing them, but not the benefit they are providing their neighbors. This results in a negative externality every time a link is severed[6]. In many situations, the

[6]The negative effects of ostracism can be particularly acute in financial networks during times of distress in which banks get shut out of funding, as is the case of a liquidity freeze.

negative effects of learning outweigh the positive effects, so on balance learning is actually harmful. In particular, increasing the learning rate about marginal agents whose neighbors are high quality is bad, because forcing the marginal quality agent out of the network sacrifices the social benefit of the link to the high quality agent. However, increasing the rate of learning about a marginal quality agent whose neighbors are also marginal quality is good, because more information will be revealed about that marginal quality agent, allowing its neighbors to more quickly sever their links to it. The impact of learning can therefore be either positive or negative depending on the specific network.

Our welfare results have important implications for network planning and are useful in a diverse range of settings, such as in guiding the formation of networks by the policies of a financial regulator, human resources department, online community, etc. Due to the varying effects of learning, we show that the optimal network design will be quite different for different groups of agents. For instance, when agents all have high initial reputations, the optimal network allows all agents to be fully connected (so that agents can benefit fully from their repeated interactions). On the other hand, if some agents have low initial reputations, then allowing all agents to connect is not optimal, and it will instead be desirable to constrain the network connections by isolating low reputation agents from each other. If such agents did link, they would both send more information about themselves through this link, causing themselves to be ostracized more quickly. Each agent, as well as the overall network, could then be worse off through the formation of this link due to the faster learning caused by the link. Therefore a star or a core-periphery network connectivity structure would generate higher social welfare than a complete network even when all agents have initial expected qualities higher than the linking cost. Such a situation arises for instance if there are two separate groups of agents, one group with

---

Ostracism has also been demonstrated in a wide variety of social settings in the social psychology literature. We discuss this literature and our model's implications in the Literature Review section.

very high reputation and the other group with moderate reputation. By placing the high reputation agents in the core and the moderate reputation agents in the periphery, the high reputation agents are able to produce large benefits for the network, and the potential harm from the moderate reputation agents is minimized[7].

Finally, we consider three extensions of our model that allow for even richer network dynamics and learning. In the first extension, we allow the mechanism designer to provide the agents with a subsidy that encourages linking[8]. The effect of such a subsidy is to promote the amount of experimentation done by the agents, and we show that a sufficiently large subsidy can always improve overall social welfare because of this. In the second extension, we allow for agents with high enough reputations to form new links with each other, and we show that social welfare will be increased when the linking threshold is high enough. In the third extension, we allow new agents to enter the network over time, and we consider the optimal time at which new agents should arrive. We show that all agents should be allowed to enter the network eventually, but delayed entry is desirable in certain networks to protect the reputations of vulnerable incumbent agents.

2. Literature Review

2.1. Relation to Theoretical Networks Literature

Our paper represents a novel contribution to the network formation literature, by being among the first to consider incomplete information and learning in networks, as well as by providing a tractable model that allows for the computation of many properties, including the *ex ante* social welfare, of different

[7]This provides a new reputational reason for the benefits of a core-periphery network, in contrast to other, non-informational, reasons that have been proposed in the networks literature.

[8]For instance, a financial regulator could guarantee transactions within a financial network to make them less risky.

network topologies. Other papers in the network literature have usually studied network dynamics only in settings of complete information where agents perfectly know each other's qualities. For example, the papers [4], [5], [6], and [7] all consider networks where the agents have complete information. In these models, agents are aware of the exact qualities of all other agents and there is no learning. The network dynamics arise instead from externalities and indirect benefits between agents that are not directly linked. For some networks, such as communication networks, these indirect benefits seem important, as an agent who has many high quality neighbors will likely be able to transmit higher quality information as well. However, in other networks such as friendship networks these indirect benefits are less relevant and it is the specific quality of each individual agent that is most important. This is especially applicable in situations where a new group of agents are meeting for the first time and learning about each other through mutual interactions.

We do not assume any indirect benefits in our model and focus instead on the dynamics resulting from incomplete information and learning. Agent learning strongly influences the network formation process in a way that would not arise with complete information. Agents that send good signals will develop high reputations and remain in the network, whereas agents that send bad signals will develop low reputations and eventually become ostracized by having their neighbors cut off links. The rate of learning about an agent's quality affects how quickly the network evolves and has a strong impact on the resulting social welfare. With complete information however, such dynamics would not occur because agents would know each other's qualities perfectly at the onset. For instance, [6] considers a dynamic network formation model where agents form links under complete information. When there are no indirect benefits between agents in that paper's model, each agent would make a one time linking decision with every other agent and never update its choice later on. But with learning agents may change their linking choices by breaking off links with neighbors that consistently produce low benefits. Incomplete information causes links to fluctuate dynamically over time as new information arrives and beliefs are updated, instead of staying static as in the complete information case. We assert that such dynamics are key and even the main driver of network evolution among a group of agents that is uncertain of each other's type.

In addition, the tractability of our model allows us to explicitly compute the social welfare for different network structures even under incomplete information. This tractability arises from the use of continuous time diffusion processes in our model, which allows for closed form equations of the probabilities that different networks emerge. In contrast other networks papers such as [4] and [5] use discrete time models that do not allow for such clean closed form expressions. While these other papers analyze the efficiency properties of a given fixed network, our welfare results are much stronger and allow the network to evolve endogenously over time as agents learn and update their linking decisions. This enables us to compare the *ex ante* optimality of different initial network structures, as well as provide general results for when specific network structures are optimal. For instance, we show that when the rate of learning in the network is either very slow or very fast, a complete network is optimal if the agent's initial expected qualities are all higher than the cost of maintaining a link. But when learning is at an intermediate rate, it may be optimal to prevent vulnerable agents from connecting, even if their initial expected qualities are higher than the linking cost, due to the negative externalities associated with reputational effects. Such a result cannot arise under complete information, where if agent's qualities are all perfectly known it would be strictly better for all of them to be linked initially.

This paper is also tied to the literature on observational learning in networks, such as [8], [9], and [10]. In these models there is a fixed exogenous network on which the agents interact, and the agents learn about an exogenous state of the world through this network by observing the actions of neighbors. These papers provide results regarding the speed and accuracy of the observational learning that can be achieved by agents connected through different types of networks. Our paper is significantly different from this literature because agents learn about other agents' *qualities* instead of an *exogenous* state of the world. As such, agents will wish to update their linking decisions over time as their beliefs about the agents with whom they are connected with change. The network and learning *co-evolve* in our model, causing the network structure to evolve *endogenously*.

Two related papers that feature information in networks are [11] and [12]. [11] focuses on the issue of moral hazard and monitoring, and it considers the diffusion of information about agent actions across a network. The paper analyzes how the structures of the networks that emerge is affected by the transmission of information about player actions, and it shows through simulations and mean-field analysis that the inclusion of network-based information can increase network density. Our work instead focuses on the issue of adverse selection and learning about agent types. We show that more information can be lower network density because it leads to greater ostracization among agents. [12] also considers learning by agents about the types of other agents within a network, and it shows how incomplete information and the learning process can lead to a wide variety of network structures and dynamics. However, this paper considers a discrete time model and incorporates a simplified learning process in which information is revealed *immediately*, after a single interaction. On the contrary, in our model information is revealed gradually, and the linking decisions and learning occur simultaneously. We are thus able to analyze how the precise rate of learning and its interaction with the network structure affects network dynamics and social welfare. Furthermore, the tractability of our model compared to both papers allows us to consider the social welfare generated across the entire path of network evolution, as opposed to the welfare of the long run average network. We are therefore able to address issues of network design from an *ex ante* perspective.

Finally, another of our papers [13] (forthcoming) considers agents that learn via private signals about others. The paper also assumes that agent types are distributed on a line, and agents prefer to link with other agents that have similar types, that is they have homophilic preferences for forming links. Thus the type of learning and the preferences that agents have are both different from the current paper. We show in the other paper that incomplete information leads to the formation of sparser networks, and that the difference between complete and incomplete information networks is largest when the level of homophily is intermediate.

2.2. Relation to Financial Networks Literature

Our paper is also related to the growing financial networks literature. There have been numerous recent papers which seek to explain the prevalent core-periphery structure of financial networks. Such core-periphery financial network structures have been well documented empirically in a variety of markets, such as for municipal bonds ([14]) and securitization ([15]). The theoretical papers of [16], [17], [18], [19], and [20] all propose models that seek to explain the prevalence of core-periphery networks. These papers show that features such as various forms of dealer heterogeneity can result in core-periphery type structures.

However most of these papers operate in complete information settings where the types of other agents are directly observable. The relatively few papers that do consider incomplete information focus instead on learning through investment in information gathering (regarding debt repayment) rather than on learning through interactions which are affected by the network structure itself. For instance, [19] shows that since star networks can allow for efficient mutual monitoring by financial institutions, they also lead to more efficient trading. [21] shows that the benefits that a core-periphery network provides leads to greater stability over time.

Our paper also provides a justification for the multitude of real world core-periphery networks, as we show that such core-periphery networks maximize social welfare in certain networks where agents vary in reputation. However, our result is driven by the presence of reputational forces, unlike the previous papers. In our model, a core-periphery network lowers the reputational risks for vulnerable low-reputation agents, and can thus prevent them from being shut out of the financial network as quickly.

Furthermore the setting of our paper is different from the settings of the other papers. The papers that consider complete information are more relevant for longer time frames and stable financial market conditions where informational uncertainty about counterparties is low. We view our model instead as describing a short time period with great uncertainty. For instance in the aftermath of a financial crisis, banks are very unsure of the solvency of other banks due to the difficulty of assessing the quality of their assets. In such situations, banks will be hesitant to trade with each other and will carefully attempt to learn the solvency of other institutions through observations of repayments. Thus each bank's reputation evolves over time. Banks that obtain low reputations may get shut out of the funding market entirely during liquidity runs, as was the case during the collapse of Lehman Brothers in the 2007-2008 financial crisis. It is important for a financial regulator to carefully structure the trading network and control the interactions so that such situations can be mitigated.

2.3. Relation to Social Ostracism Literature

Finally, we note that our model also has important implications for social and organizational networks. Our results about the negative externalities of reputational learning highlight the damaging impacts of ostracism found in the social psychology literature. Social ostracism is a prevalent force that has been well documented in the social psychology literature in numerous settings ranging from online interactions to office workplaces. As [22] states, "Social ostracism is a pervasive and ubiquitous phenomenon." In this literature, ostracism can also occur when an agent's perceived quality drops too low, and has harmful effects on the ostracized agent. As the paper by [23] notes, "Ostracism is a common, yet painful social experience...Individuals who do not fit the group's definition of a contributing member may find themselves a likely candidate for *punitive* ostracism". That paper shows the occurrence of ostracism via an online experiment where agents differ in their ability to play a game. Agents who play badly became ostracized by the others. This effect is similar to our model, where agents who are learned to be of low quality are ostracized from the network.

Ostracism can also occur in workplaces, as employees may be ostracized by their coworkers. [24] notes that "not only are such experiences extremely painful, but under some circumstances they can have an even greater negative impact than other harmful workplace behaviors such as aggression and harassment." It is therefore important for companies to consider the harmful effects of ostracism that can occur through workplace interactions. We provide guidelines for minimizing the negative effects of ostracism through placing lower reputation agents in less central positions of the network.

3. Model

3.1. Overview

We consider an infinite horizon continuous time model with a finite set of agents denoted by $V = \{1, 2, ..., N\}$. At every moment in time, the agents choose which other agents to link with, and a link is established only under mutual consent. These choices are made subject to an underlying network constraint $\Omega = \{\omega_{ij}\}$ that specifies the pairs of agents that are able to link with each other[9]. For each pair of agents, $\omega_{ij} = 1$ if agents i and j can connect with each other and $\omega_{ij} = 0$ otherwise. We call agents i and j neighbors if they can connect. Initially (time $t = 0$), agents are linked according to a network $G^0 = \{g_{ij}^0\} \subseteq \Omega$. As the network will change over time, we denote G^t as the network at time t. Moreover, we let $k_i^t = \sum_j g_{ij}^t$ be the number of links that agent i has at time t, and we let K_i^t denote the set of neighbors of agent i at time t.

Agents receive flow payoffs from each link equal to the benefit of that link minus the cost. Each agent i must pay a flow cost c for each of its links that is

[9]This network constraint Ω may arise from the specific interests/desires of the agents regarding who they want to link with, or from potential physical/geographical constraints that limit agents from linking. It may also be planned, e.g. through the policies of a financial regulator for a network of financial institutions, or by the human resources department in a company for a network of employees.

active. Hence, at time t, agent i pays a total cost of $k_i^t c$ for all its links. Agents also obtain benefits from their links, depending on their linked neighbors' qualities q_i. However each agent's true quality is initially unknown to all agents, and we do not require that agents know their own qualities. At the start of the model, each agent i's quality q_i is drawn from a commonly known normal distribution $\mathcal{N}(\mu_i, \sigma_i^2)$ with $\mu_i > c$. Both the mean and the variance are allowed to vary across agents, and several of our results below will utilize this heterogeneity. Agent i generates a different noisy benefit $b_{ij}(t)$ for each agent j that is linked to it, and these benefits follow a diffusion $db_{ij}(t) = q_i dt + \tau_i^{-1/2} dZ_{ij}(t)$, where the drift is the true quality q_i and the variance depends on τ_i, an exogenous parameter we call the *signal precision* of agent i[10]. $Z_{ij}(t)$ is a standard zero-drift and unit variance Brownian motion and represents the random fluctuations in the benefits of each interaction. $Z_{ij}(t)$ is assumed to be independent over all i and j, and therefore all the benefits produced by agent i are conditionally independent given q_i. We assume that all the benefits that agent i produces are observed by all the neighbors of i, which ensures that agent i's neighbors all have the same beliefs about i at any point in time (information is locally public among agent i's neighbors)[11]. For each agent i, we define the agent's benefit history as the history of all previous benefits, $\mathcal{H}_i^t = \{b_{ij}^{t'}\}_{t'=0}^t$.

We assume that agents are myopic, and they thus consider only the current

[10]We can think of the signal precision as representing how much information the agent reveals about itself in each interaction, with a higher precision corresponding to more information. It could depend on the type of interaction with the agent (e.g. close partnerships vs. chance encounters), or factors like the agent's personality.

[11]This is an important assumption to maintain the tractability of the model. One possible interpretation is an online expertise network where the output of agent i is public, so that all neighbors of agent i can judge the benefit that i has provided to all its links. Or in an offline setting, we could assume that the neighbors of agent i are continuously discussing the benefits they have received from i with all the other neighbors of i, so that the neighbors maintain the same beliefs. For most of our results, the information does not need to be fully public; the information regarding agent i needs only be available to all the direct neighbors of agent i.

flow benefit when making linking decisions[12]. Each agent's utility is assumed to be linear in the benefits provided by each link and the linking cost. This also implies that agents are risk neutral and so consider the expectation over neighbor qualities when there is uncertainty. The flow utility of agent i at any time t is given by the following equation:

$$U_i = \sum_{\{j \in K_i^t\}} (E[q_j | \mathcal{H}_i^t] - c) \tag{1}$$

3.2. Reputation and Learning Speed

Since we have assumed a diffusion process, a sufficient statistic for all the individual link benefits is the average benefit per link produced by agent i up to time t, which we denote as $B_i(t)$. Given our above assumptions, $B_i(t)$ follows a diffusion $dB_i(t) = q_i dt + (k_i^t \tau_i)^{-1/2} dZ_i(t)$ where the drift rate is the true quality q_i, the instantaneous volatility rate $(k_i^t \tau_i)^{-1/2}$ depends on the number of links agent i has at time t, and $Z_i(t)$ is the standard Brownian motion with zero-drift and unit-variance. Importantly, this equation shows that the more links an agent has, the lower its volatility rate and the faster its true quality q_i is learned. This is because an agent with more links produces more individual benefits, and so the average over all benefits is more precise. Note also that an agent with no links would not send any information, and thus there would be no learning about its quality. Therefore the topology of the network determines the rate of learning about each agent's quality.

If at time t all the links of agent i are severed, then no benefit will be produced by agent i and this will be denoted as $b_i^t = \emptyset$. In this case no information is added, and thus the diffusion of agent i is stopped at its current level. As mentioned, the prior belief of agent i's quality is $\mathcal{N}(\mu_i, \sigma_i^2)$, and agents will up-

[12] Such an assumption is common within the networks literature to maintain tractability, see [4] or [6] for instance. Myopia is an appropriate assumption in financial networks where firm managers have myopic incentives. Such myopic incentives have been documented empirically in papers such as [25] and [26]. We relax this assumption in the extensions section where we allow for subsidies that change agent linking strategies.

date this belief in a Bayesian fashion in light of the observations of flow benefits. These observations combined with the prior quality distribution will result in a posterior belief distribution of agent i's quality $f(q_i|\mathcal{H}_i^t)$ which is also normally distributed[13]. We denote $\mu_i^t = E[q_i|\mathcal{H}_i^t]$ as the expected quality of agent i given the history $\mathcal{H}_i^t$ and call it the *reputation* of agent i at time t. The reputation represents the expected flow benefit of linking with agent i at time t.

We have assumed that agents are myopic. Therefore, to maximize flow utilities, agent i will cut off its link with agent j once agent j's reputation μ_j^t falls below the linking cost c. Since we assume all agents have homogeneous linking costs, and all neighbors have the same beliefs, any other agent that is linked to j will also decide to sever its link. From this moment on, agent j is effectively ostracized from the network; since it no longer has any links it cannot send any further information that could potentially improve its reputation[14]. While in the base model an ostracized agent cannot return to the network, we relax this assumption in the extensions section.

4. Network Dynamics and Stability

4.1. Network Dynamics

The dynamics of the model evolve as follows: all pairs of agents that are neighbors according to the network constraint Ω will choose to link at time zero, since we have assumed that all agents have initial reputations higher than the cost c (any agent with an initial reputation lower than c is immediately ostracized from the network and would not need to be considered). Therefore the initial network at time 0 will be the same as the network constraint, $G^0 = \Omega$.

[13] As mentioned a sufficient statistic for the entire history is $B_i(t)$, so a neighbor only needs to know $B_i(t)$ in order to calculate this posterior.

[14] Although ostracism may seem harsh, as we noted earlier ostracism is a prevelant phenomenon that has been widely studied in the social psychology literature, in settings ranging from online interactions to office workplaces. Furthermore, in financial networks low reputation institutions may get shut out of funding completely during liquidity crisis.

Over time agents that send bad signals will have their reputations decrease, and once an agent's reputation hits c its neighbors will no longer wish to link with it. All its neighbors will sever their links and the agent is effectively ostracized from the network. We will show that this always happens for an agent with true quality $q_i \leq c$, and will still happen with positive probability for an agent with quality $q_i > c$. The ostracization of an agent will affect its former neighbors as well. Since they now have once less link each, they will produce information about themselves more slowly than before, and so their reputations will be updated less quickly.

The remaining agents in the network will continue to link and send signals until someone else's reputation drops too low and that agent is also ostracized. This process will continue until the qualities of all the remaining agents are known with very high precision and in the limit their reputations no longer change. Since agent qualities are fixed, by the law of large numbers any agents that remain in the network will have their qualities learned perfectly in the limit as $t \to \infty$, and the network will tend towards a limiting structure that we call the *stable network*. The next section will explicitly characterize these stable networks, but we note that many different stable networks could potentially emerge depending on the true qualities of the agents and the signals they produce.

4.2. Stable Networks

As mentioned, we call the limiting network structure as t goes to infinity, denoted by G^∞, a stable network. Formally, let $G^\infty \equiv \lim_{t\to\infty} G^t$. This limiting structure always exists since agent qualities are fixed, so by the law of large numbers any agent that remains in the network will have its quality learned to an arbitrary precision over time. The probability that an agent who is still in the network at time t ever becomes ostracized must therefore tend to zero as $t \to \infty$ (we show this analytically below). Which specific stable network eventually emerges is random and depends on the signal realizations of each agent. The tractability of our model allows us to explicitly characterize the set of stable networks that could emerge given a set of agents and a network constraint Ω, as well as the impact of the rate of learning on the probability distribution over stable networks.

To understand which stable networks G^{∞} can emerge, we investigate whether a link l_{ij} between agents i, j can exist at $t = \infty$. If two agents i and j are not neighbors (i.e. $\omega_{ij} = 0$), then it is certain that $g_{ij}^{\infty} = 0$. If two agents i and j are neighbors (i.e. $\omega_{ij} = 1$), then the existence of this link l_{ij} at $t = \infty$ requires that the reputations of both i and j never hit c for all finite t, which means that neither agent is ever ostracized. Hence G^{∞} will always be a subset of the initial network G^0, and is composed only of agents whose reputations never hit c for all finite t.

We say that an agent is *included in the stable network* if their reputation never hits c for all t, so that they are never ostracized from the network[15]. Note that being included in the stable network does not imply that an agent has any links in the stable network, as it could also be that all of the neighbors of that agent were ostracized even though the agent itself was not. We can calculate the *ex ante* probability that an agent i is included in the stable network, which we denote by $P(S_i)$ with S_i denoting the event in which agent i is included in the stable network. This can be accomplished using standard results regarding Brownian motion hitting probabilities, since $P(S_i)$ is equal to the probability that the agent's reputation never hits c for all finite t. The following proposition gives this probability. All proofs will be provided in the appendix.

Proposition 1. $P(S_i)$ *depends only on the initial quality distribution and the*

[15] As a technical note, when we make the ostracization classification, we assume that an agent who has all its neighbors ostracized continues to send information about itself at its signal precision level, with the signals sent via the same probability distribution which is based on its true quality. So we still considered the agent "ostracized" if its reputation drops to c via this information process even after all its neighbors have been ostracized. This assumption is made for technical purposes only and has no impact on the dynamics or the welfare of the model, as the agent has no links in this case.

link cost and can be computed by

$$P(S_i) = \int_c^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i/\sigma_i \qquad (2)$$

Proposition 1 has several important implications. Note that since $P(S_i)$ is positive and less than 1 for all i, no agent is certain to be included in or excluded from the stable network. Also note that the probability an agent is part of the stable network is independent of that agent's signal precision τ_i. Therefore the rate at which the agent sends information does not affect the chance that it is in the stable network. This is because the rate at which the agent sends information only affects *when* it gets ostracized from the network, but not *if* it gets ostracized overall[16]. Furthermore, note that the probability an agent i is included in the stable network is independent of its links with other agents and the properties of those agents. Connections with other agents affect the rate at which an agent sends information but not the agent's true quality, and so will not impact whether it is eventually ostracized from the network.

Using the explicit expression above, we can also describe how $P(S_i)$ depends on an agent's initial mean and variance, μ_i and σ_i.

Corollary 1. *For each agent i, $P(S_i)$ is increasing in the mean of its initial quality μ_i, decreasing in the variance of its initial quality σ_i^2, and decreasing in the link cost c. Moreover, $\lim_{\mu_i \to \infty} P(S_i) = 1$, $\lim_{\sigma_i \to 0} P(S_i) = 1$, $\lim_{c \to -\infty} P(S_i) = 1$.*

These properties are intuitive since an agent with a higher mean quality and lower variance is less likely to have its reputation drop below c, and so is less

[16]Intuitively, recall that reputation evolves through Bayes updating of the Brownian motion. A higher precision increases the amount of information sent at every moment in time, but the overall probability distribution of the information that is sent across all time remains the same. This can also be shown rigorously using the formula for the survival probability of an agent in the proof of Proposition 1 in the appendix, which depends on τ_i only through the term $t\tau_i$. Therefore increasing τ_i and decreasing the considered time t proportionally leaves the overall survival probability unchanged.

likely to become ostracized. Moreover, lowering the linking cost also reduces the hitting probability since the agent's reputation would now have to fall further to be excluded from the network.

As mentioned, G^{∞} must be a subset of G^0. Further, it can only contain links among pairs of agents that are both included in the stable network and linked in the initial network. Equivalently, the set of stable networks can be thought of as the set of networks that can be reached from G^0 by sequentially ostracizing agents. Let $I\{S_i\}$ denote the indicator variable of the event in which agent i is included in the stable network. Formally, a network can be stable if and only if it is a matrix with entries given by $g_{ij} = I\{S_i\}I\{S_j\}I\{g_{ij}^0 = 1\}$, for some realization of $\{S_i, ^\neg S_i\}_{i \in V}$. Links can exist only among agents that were never ostracized and were linked in the original network. Note that different realizations of $\{S_i, ^\neg S_i\}_{i \in V}$ could potentially correspond to the same stable network[17].

By Proposition 1, we know that the rates of learning do not affect the probability of each event S_i. Since the rate of learning has no effect at an individual level, it cannot have an effect at the aggregate level either. This is formalized in the next theorem. We can also use the equation in Proposition 1 to derive an analytic expression for the probability that any specific stable network emerges, which is presented in the corollary below. Figure A.1 in the appendix shows an example of how the corollary can be applied to a simple network of three agents.

Theorem 1. *The signal precisions of the agents, $\{\tau_i\}_{i \in V}$, do not affect the set of stable networks that can emerge or the probability that any stable network emerges.*

Corollary 2. *The probability that a network G is a stable network is given by*

[17]For instance suppose that the network comprises only two agents i and j. Then the event in which S_i but not S_j occurs and the event in which both S_i and S_j occur lead to the same stable network structure: the empty network.

$\sum_{\{S_i\}} \prod_i P(S_i)$ *where the summation is over all realizations of* $\{S_i, \neg S_i\}_{i \in V}$ *that correspond to* G.

We have shown that the speed of learning has no impact on the probability that a network G is stable. This is intuitive since learning only affects the duration of a link but not its final state. However, learning will have a crucial role on the social welfare of a network, which directly depends on how long the agents are connected. We will consider the impact of learning on the social welfare in the next section.

5. Welfare Computation

We will analyze overall social welfare from an *ex ante* perspective, given only the network constraint Ω and the prior agent quality distributions. Importantly the *ex ante* welfare is calculated before the agent qualities are learned and any signals are sent. This type of welfare is the most suitable for the type of design settings we will consider later, as it requires the least knowledge on the part of the network designer. Let $P(L_{ij}^t|q, G^0)$ denote the probability that the link between agents i and j still exists at time t. Also, let the parameter ρ represent the discount rate of the network designer[18]. We can define the overall *ex ante* social welfare W formally as follows:

$$W = \int_{q_1=-\infty}^{\infty} ... \int_{q_N=-\infty}^{\infty} \sum_{i,j} \int_0^{\infty} e^{-\rho t}(q_j - c) P(L_{ij}^t|q, G^0) dt \phi(\frac{q_N - \mu_N}{\sigma_N}) dq_N/\sigma_N ... \phi(\frac{q_1 - \mu_1}{\sigma_1}) dq_1/\sigma_1$$

We will show that this social welfare expression can be calculated in a tractable fashion using a somewhat indirect approach. This approach utilizes the fact that the *ex ante* social welfare is an expectation over all the possible *ex post* signal realizations. We can calculate the *ex ante* welfare by integrating over all possible realizations of the *ex post* welfare, which simplifies the equation for *ex ante* social welfare to a much more tractable form.

[18] We are assuming that the designer itself is more patient than the myopic agents. This can be interpreted, for instance, as a company manager who is more patient than its workers who act myopically in their interactions, or a financial regulator that is more patient than the financial institutions, which have managers with myopic incentives.

5.1. Ex post welfare

Consider an *ex post* realization of agent hitting times $\varepsilon = \{\varepsilon_i^{t_i}\}_{i\in\mathcal{V}}$, where $\varepsilon_i^{t_i}$ denotes the event in which agent i's reputation hits c at time t_i given all the agent signals (note that $t_i = \infty$ means that agent i's reputation never hits c). In the event in which $t_i < \infty$, since the belief at time t_i is correct, the expected value of agent i's quality conditional on this event $\varepsilon_i^{t_i}$ is $E[q_i|\varepsilon_i^{t_i<\infty}] = c$. In the event with $t_i = \infty$, since the initial belief is accurate in expectation

$$\mu_i = E[q_i] = P(\varepsilon_i^{t_i<\infty})E[q_i|\varepsilon_i^{t_i<\infty}] + P(\varepsilon_i^{t_i=\infty})E[q_i|\varepsilon_i^{t_i=\infty}] \quad (3)$$

$$= (1 - P(S_i))c + P(S_i)E[q_i|\varepsilon_i^{t_i=\infty}] \quad (4)$$

and we have

$$E[q_i|\varepsilon_i^{t_i=\infty}] = \frac{\mu_i - (1 - P(S_i))c}{P(S_i)} \quad (5)$$

where $P(S_i)$ is given by Proposition 1 and is independent of the network and the learning speed.

According to the above discussion, given an *ex post* realization ε, an agent i obtains 0 surplus from its neighbors that have finite hitting times and obtains positive surplus from those whose reputation never hits c (and are therefore included in the stable network). The exact benefit agent i receives in the second case depends on its own hitting time t_i, which determines the link breaking time with the other agents. We can calculate the *ex post* surplus that an agent i receives given ε as follows:

$$W_i(\varepsilon) = E_{q|\varepsilon}\left[\sum_{j:g_{ij}^0=1}\int_0^{\min\{t_i,t_j\}} e^{-\rho t}(q_j - c)dt\right] \quad (6)$$

$$= \sum_{j:g_{ij}^0=1,t_j=\infty}\int_0^{t_i} e^{-\rho t}\left(\frac{\mu_j - (1 - P(S_j))c}{P(S_j)} - c\right)dt \quad (7)$$

$$= \frac{1 - e^{-\rho t_i}}{\rho}\sum_{j:g_{ij}^0=1,t_j=\infty}\frac{\mu_j - c}{P(S_j)} \quad (8)$$

Note that this W_i is taken from the perspective of the designer as it incorporates futures payoffs at the discount rate of ρ. This equation shows that in each *ex* *post* realization of other agent hitting times, agent i benefits if t_i increases and it is ostracized later from the network. Summing over all agents, the social welfare given the *ex post* realization ε is therefore

$$W(\varepsilon) = \sum_i \left(\frac{1 - e^{-\rho t_i}}{\rho} \sum_{j: g_{ij}^0 = 1, t_j = \infty} \frac{\mu_j - c}{P(S_j)} \right) \tag{9}$$

By taking the expectation over the events ε, the *ex ante* social welfare can be given by $W = E_\varepsilon[W(\varepsilon)]$. In order to compute the *ex ante* social welfare, we still need to know the distribution of the t_i, which is coupled in a complicated manner with the initial network and the learning process. For instance, if the neighbor of agent i has a low hitting time and is ostracized quickly, then agent i sends information at a slower rate and its own hitting time would increase. Thus directly computing the social welfare using the above equation is still difficult. In the next subsection, we develop an indirect method to calculate the distribution of t_i.

5.2. Hitting time mapping

Recall that an agent's links will scale up the rate at which it sends information compared to the rate it would send information if its precision were constant at the base level of τ_i. Therefore each link also scales down the time at which the agent's reputation hits c. So to calculate when the agent is ostracized, we can first find when the agent's reputation would hit c through sending signals at its signal precision level, and then scale this time downwards proportionately based on the network effect[19]. Consider an *ex post* realization of hitting times $\hat{\varepsilon} = \{\hat{\varepsilon}_i^{t_i}\}_{i \in \mathcal{V}}$ in which agent i's reputation would hit c at time t_i if its precision were fixed at τ_i at all times. Note that the events $\hat{\varepsilon}_i^{t_i}$ are independent from each other across different agents, and since the precision is fixed they also do not depend on the network structure. The probability of $\hat{\varepsilon}_i^{t_i}$ is explicitly computed in the following lemma.

[19]Refer to footnote 16 for a justification of this type of scaling.

Lemma 1. *The probability density function $f(\hat{\varepsilon}_i^{t_i}), \forall t_i < \infty$ can be computed as*

$$f(\hat{\varepsilon}_i^{t_i}) = \int_{-\infty}^{\infty} \frac{\mu_i - c}{\sigma_i^2 \sqrt{\tau_i}} t^{-3/2} \phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \phi(\frac{q_i - \mu_i}{\sigma}) dq_i / \sigma_i \tag{10}$$

The probability mass point function $f(\hat{\varepsilon}_i^{t_i=\infty}) = P(S_i)$.

Using Lemma 1, we can easily obtain the distribution of joint events $f(\hat{\varepsilon}) = \prod_i f(\hat{\varepsilon}_i^{t_i})$ due to the fact that the individual events are independent. This would measure the joint probability of the agents exiting the network at times $\{t_i\}_{i \in \mathcal{V}}$ if the information sending speed of the agents were not being scaled by the number of their links. If there were no network effect, the *ex ante* social welfare could be directly computed using the distribution of hitting times given by Lemma 1. However, due to the network effect, the actual hitting times may vary for each $\hat{\varepsilon}$. We can define $M : [0, \infty]^N \to [0, \infty]^N$ to be the hitting time mapping function, which maps the hitting times with no network effect to the actual hitting times when there is a network effect. In the appendix in Figure A.2 we present an algorithm for computing M, which operates by scaling the information speed of each agent at every time t by their current number of neighbors and updating the speed at which an agent sends information when a neighbor is ostracized. Note that if $t_i = \infty$ in the event $\hat{\varepsilon}_i^{t_i}$ then it is also ∞ in the mapped event $\varepsilon_i^{t_i}$. This means that an agent that never leaves the network with no scaling effect will not leave when the times are scaled either. Then given a realization $\hat{\varepsilon}$, the *ex post* social surplus can be computed as

$$W(\hat{\varepsilon}) = \sum_i \left(\frac{1 - e^{-\rho M_i(t)}}{\rho} \sum_{j: g_{ij}^0 = 1, t_j = \infty} \frac{\mu_j - c}{P(S_j)} \right) \tag{11}$$

Therefore, the *ex ante* social welfare is $W = E_{\hat{\varepsilon}}[W(\hat{\varepsilon})]$. We note that this is a tractable equation for the *ex ante* social welfare given any network structure and set of agents. Proposition 1 gives the explicit expression for $P(S_j)$, and Lemma 1 provides the distribution of $\hat{\varepsilon}$. Thus our model allows for easy and tractable computations of the *ex ante* social welfare of any type of network. Theorem 2 below formalizes this result.

Theorem 2. *Given Ω, the initial quality distributions, and the link cost c, the overall ex ante social welfare can be computed as follows*

$$W = E_{\hat{\varepsilon}}\left[\sum_i \left(\frac{1 - e^{-\rho M_i(t)}}{\rho} \sum_{j:g_{ij}^0=1, t_j=\infty} \frac{\mu_j - c}{P(S_j)}\right)\right] \tag{12}$$

where the distribution of $\hat{\varepsilon}$ is computed using Lemma 1 and the hitting time mapping function M is given in the appendix.

6. Impact of Information and Learning

In this section we study the impact of learning on *ex ante* welfare, both individual and overall, given an initial network G^0. In particular, we will show how the agents' signal precisions, a representation of the rate of learning, impact individual agent welfare as well as the overall social welfare.

As a benchmark, we consider the social welfare when there is no learning, which we denote by W^*. When there is no learning, no existing link will be severed. The social welfare of an agent i without learning can therefore be computed by summing over the mean qualities of all agents it is connected with initially:

$$W_i^* = \sum_{j:g_{ij}^0=1} \int_0^\infty e^{-\rho t}(\mu_j - c)dt = \frac{1}{\rho} \sum_{j:g_{ij}^0=1} (\mu_j - c) \tag{13}$$

The *ex ante* overall social welfare without learning is given by the sum over the individual welfares:

$$W^* = \sum_i W_i^* = \frac{1}{\rho} \sum_i \sum_{j:g_{ij}^0=1} (\mu_j - c) \tag{14}$$

6.1. Overall Impact of Learning

Let $W(\tau_1, ..., \tau_N)$ be the *ex ante* social welfare when agents learn each other's true quality with the signal precisions being $\tau_1, ..., \tau_N$. We also let $W_i(\tau_1, ..., \tau_N)$ represent an agent i's *ex ante* welfare given these signal precisions. The next theorem states that in any network, the addition of learning has a negative impact on every individual's *ex ante* welfare for any value of the signal precisions. This immediately implies that it lowers the overall *ex ante* social welfare as well.

Theorem 3. *$W_i(\tau_1, ..., \tau_N) < W_i^*$ for all i and for all $\tau_1, ...\tau_N$.*

There are two main factors that are at work in this result. First, the myopia of the agents causes the learning to be done inefficiently. Second, cutting off a link imposes a negative externality on the agent who is ostracized, since that agent can no longer receive benefits from its neighbors. Taken together, these factors lead to a reduction in overall social welfare. More precisely, when a link l_{ij} is severed due to agent j's reputation hitting c, agent i does not gain welfare compared to the case without learning. This is because the expected value of having a link with i from t_j^* on is 0 and thus having the link or not makes no difference[20]. However, agent j loses welfare compared to the case without learning because agent i's reputation is still above the link cost and thus having the link would benefit j over not having the link.

This result supports the damaging impacts of ostracism found in the social psychology literature, which were mentioned above in the literature review. The social psychology literature usually documents the harmful effects of ostracism from the perspective of the agents that have become ostracized and can no longer benefit from interactions with the other agents. However, our result goes further by stating that the possibility of ostracism will actually lower *every* agent's social welfare from an *ex ante* perspective. By allowing for the ostracism of others, agents open themselves up to ostracism as well, which lowers their own welfare by more than they benefit from ostracizing other agents. Theorem 3 shows that every agent is hurt *ex ante* by ostracism, even those that wouldn't themselves be ostracized in the majority of the *ex post* realizations of the network.

[20] Agent myopia is causing the cut-off value to be too high, and so the agent does not benefit from its learning. This feature of reputational learning is similar to that shown in [27]. In Section VIII we discuss a possible solution for this problem by providing agents a subsidy to increase experimentation.

6.2. Impact of Individual Information

The previous result showed that learning is harmful on aggregate: under learning both individual and overall network welfare are lower than without learning. However, we show in this subsection that learning need not be harmful at an individual level, as the rate that a single agent sends information changes. We now investigate more closely how the information generation rate of a single agent (i.e. an agent's signal precision) affects welfare. The faster an agent generates information about its own reputation, the faster the other agents will learn its true quality (if the link is not broken).

First we characterize the impact of an agent's signal precision on that agent's own welfare. The next proposition shows that sending more information about itself will always harm an agent.

Proposition 2. *$W_i(\tau_i, \tau_{-i})$ is strictly decreasing in τ_i.*

This result is in accordance with Theorem 3 and shows that an agent sending information about itself will strictly decrease its own welfare. This is because in each realization in which the agent is ostracized from the network, the agent will now be ostracized sooner and hence it will enjoy less benefit from others. Since the agent already starts out with the maximal amount of links it can obtain, it in effect has nothing to gain and everything to lose by allowing its own reputation to vary. We relax this assumption in the extensions section and allow agents to form new links with those they are not connected with initially; under those circumstances an agent will be able to benefit by generating more information about itself.

Though increasing the information sending speed is always harmful for an agent itself, it can actually be helpful to its direct neighbors. The next proposition provides a sufficient condition on the initial network such that this holds.

Proposition 3. *Given an initial network G^0, for any two initially linked agents i and j that are linked only through a unique path (i.e. their direct link), increasing one's precision increases the other's welfare.*

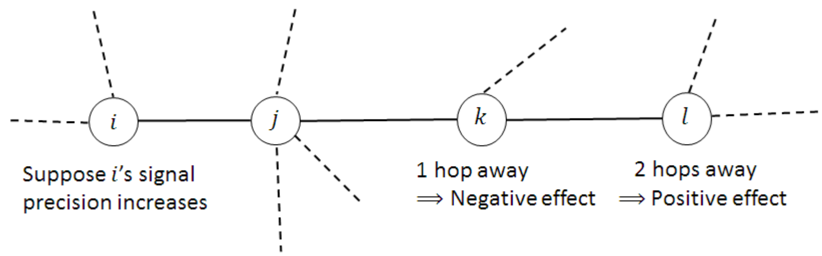

Figure 1: Example for Corollary 3

Since the information sending speed of agent j slows after agent i is ostracized, agent j's hitting time is larger. Agent j therefore prefers its direct neighbor to send more information, so that it can cut off more quickly in case the neighbor is bad. After the link is broken, agent j will also be able to reveal less information about itself, which is beneficial according to Proposition 2. In this way agent j would enjoy more benefits for a longer time from its links with its other neighbors. We can extend this analysis for more distant agents when the two agents are connected through a unique path. This is summarized in the corollary to Proposition 3 below.

Corollary 3. *Given any initial network G^0, for any two agents i and j that have a unique path between them, increasing one's signal precision decreases/increases the other's welfare if they are an odd/even number of hops away from each other.*

The above result shows an odd-even effect of the distance between two agents on the agent's welfare. In all minimally connected networks (such as star, tree, forest networks), any two agents have a unique path between each other and so the impact of any agent's information sending speed on any other agent's welfare can be completely characterized.

As an example, consider a network where four agents i, j, k, l are connected via a unique path, as depicted in Figure 1. Agent i is linked with agent j, agent j is linked with agent k, and agent k is linked with agent l. Then if agent i sends more information about itself, it stays connected with agent j for a shorter period of time. This causes agent j to send less information about itself, causing agent k to cut off its link with j more slowly if j were to be ostracized. Then agent k links with its other neighbors for a shorter length of time in expectation, decreasing the *ex ante* welfare of k. Therefore agent k is hurt when the neighbor of its neighbor, agent i, sends more information. On the other hand, agent l now links with its own neighbors for a longer length of time, and so it benefits when i sends more information.

However, when there are multiple paths between agents, which implies there are cycles in the network, the impact of changing an agent's signal precision on the other agents' welfares is much less clear. The reason is that with cycles the neighbor of an agent i's neighbor may also be linked with agent i itself[21], and so the positive and negative effects of information from Corollary 3 are entangled together. The following proposition shows that even for an immediate neighbor, the impact could be totally opposite of Proposition 3 when cycles are present in the network.

Proposition 4. *If the initial network G^0 has cycles, then it is possible that increasing some agent's signal precision decreases its immediate neighbor's welfare.*

We have seen that increasing the information sending speed of an individual agent i could be both good or bad for other agents depending on their locations in the network and their relation with agent i. We note that it could similarly be good or bad for overall social welfare. So in contrast with Theorem 3, increasing the amount of information about a single agent can benefit the network overall. This would happen, for instance, if there are three agents, i, j, and k who are connected in a line, with links ij and jk. Suppose that the mean of agent k's quality is much higher than those of the other two agents. Then most of the welfare in this network comes through the link between agents j and k. If agent i sends more information, agent j would be able to preserve its link with agent k for a longer period of time, and overall social welfare would increase.

[21] This is known in the social network literature as triadic closure.

This example highlights how critical the network structure is in determining the overall impact of the level of information sent by a single agent.

7. Optimal Networks

In this section, we study which underlying network constraints Ω maximize the overall *ex ante* social welfare. Equivalently, we could think of a benevolent network planner that wishes to maximize social welfare by designing the network constraint Ω through designating which agents are able to form links with which other agents. For instance, in the financial network setting we could think of a regulator that specifies which types of financial institutions are allowed to transact with which other types of institutions in order to maximize overall social welfare[22].

7.1. Fully connected networks

One intuition is that a fully connected network, with no constraints on links, would be optimal since it results in the largest number of links initially, and we have assumed that all agents have an initial reputation higher than the linking cost c. This intuition is accurate in certain cases, for instance if the designer is extremely impatient (i.e. $\rho \to \infty$). Since the designer cares only about the initial time period, and when time is short almost no new information can be learned, it is best to design the network based on the agents' starting reputations. Surprisingly though, the fully connected network is also optimal on the other extreme, when the designer is completely patient (i.e. $\rho \to 0$). In this case, the designer cares only about the social welfare of the stable network that eventually develops, and allowing all agents to be connected initially leads to

[22] We note that many other types of objection functions are also possible instead of the overall *ex ante* social welfare. For instance the designer may wish to maximize network welfare generated over a certain time interval, or before a set deadline is reached. Or the designer may weigh the welfare of some agents more heavily than that of others. Given the tractability of our model, it may be possible to extend many of our results to these alternative settings.

the largest number of links in the final stable network. We prove these welfare results in the proposition below.

Further, note that the designer's level of patience is inversely related with the rate of learning, as faster learning means that information is revealed sooner and thus less patience is required. Therefore a similar result holds for the rate of learning: as the rate of learning becomes extremal the fully connected network becomes optimal as well. So for instance, a financial regulator should optimally let all types of financial institutions transact with each other if it is very patient or very impatient, or if rate of learning is extremely fast or slow.

Proposition 5. *1. If the designer is either completely impatient (i.e. $\rho \to \infty$) or completely patient (i.e. $\rho \to 0$), the optimal Ω is the fully connected network. 2. Fix the other parameters of the model and suppose the agents' signal precisions are all multiplied by the same constant λ. If learning becomes very fast (i.e. $\lambda \to \infty$) or very slow (i.e. $\lambda \to 0$), then the optimal Ω is the fully connected network.*

When the designer is either completely patient or impatient, the social welfare depends only on the network G^0 or G^∞, respectively. The exact hitting time does not affect the social welfare. Similarly if the learning is very slow, then the network structure always remains at G^0, and if the learning is very fast then G^∞ is realized very quickly, so in both cases a fully connected network is optimal. The idea is that in both extremes, the exact path of learning is no longer critical and so the negative externalities of information are mitigated.

For intermediate levels of patience or learning however, changes in individual agent hitting times due to linking could have a significant impact on the social welfare. We will show later that having all agents fully connected with each other is not always the optimal choice. In the next proposition though we show that the fully connected network is still optimal in the case where the agents are homogeneous and have very high initial qualities.

Proposition 6. *Suppose all agents are ex ante identical. Fixing the other parameters, there exists $\bar{\mu}$ such that if $\mu_i > \bar{\mu}$ $\forall i$, then the optimal Ω is the fully* *connected network.*

7.2. Core-periphery networks

As agents become more heterogeneous in terms of their initial expected quality, it can be optimal to constrain connections among agents. Suppose agents are divided into two separate types, and the initial mean quality of the high type agent is μ_H while the initial mean quality of the low type agent is $\mu_L < \mu_H$. We show that when the expected qualities of the two types are sufficiently different, the optimal network constraint has a core-periphery structure[23].

Theorem 4. *Suppose that there are two groups of agents, one with initial reputation μ_L and one with initial reputation μ_H. Fixing all other parameters, there exists $\bar{\mu}$ such that $\forall \mu_H > \bar{\mu}$, the optimal Ω is a core-periphery network where all high type agents are connected with all other agents and no two low type agents are connected. ($\bar{\mu}$ will depend on the other network parameters.)*

The above result shows that under the optimal network constraint, high reputation agents should be placed in the core and connected with all other agents, while low reputation agents should be placed in the periphery and not connected with other low reputation agents. Therefore agents with lower initial reputations should be placed in less central positions within the network in order to mitigate the negative effects of ostracism. Allowing low reputation agents to connect with too many other agents would increase the rate at which they send information, causing them to be ostracized sooner and hurting them more than they would gain through the direct benefits of the extra links. This core-periphery structure is commonly seen in many real-world financial networks, with large well capitalized banks in the core and smaller banks in the periphery. A reason for this could be that the greater reputation of large banks lets them

[23] Although this theorem assumes there are exactly two types, a similar result holds if instead the agents are composed of two groups and within each group they have parameters that are sufficiently close together.

withstand negative shocks more easily without being ostracized by their counterparties. Smaller banks produce less information through their lesser number of transactions, allowing them to avoid being ostracized as quickly[24].

We note that the above result depends heavily on the type of learning environment that is present. From Proposition 6, we know that if the designer was either very patient or impatient, or if learning was very slow or very fast (relative to the parameters of the agents), then the optimal initial network would be the fully connected network. Fixing any set of agent reputations, a core-periphery network structure is only optimal at intermediate levels of learning.

7.3. Star Networks

Star networks are common networks in the real world, with a single central agent is connected with many peripheral agents. Examples include a single boss and many subordinates, the head of a political party that coordinates the disparate branches of the party, or a large trader that deals with many small traders. There are several important forces to consider when placing agents within a star network. Such networks depend greatly on the central agent, because that agent is connected with all other agents and it therefore has the most links. The central agent is therefore very important, and choosing the best agent to be in the center is crucial to the overall welfare of the network.

The initial mean and the signal precision of the central agent are two exogenous parameters that must be carefully considered when choosing the central agent. A high initial mean is beneficial because it increases the expected flow benefits that all the other agents who are connected to the central agent will

[24] We note that financial regulators have started imposing core-periphery structures on various financial networks to encourage stability. Many standardized OTC derivatives are now required to be cleared through a central clearing counterparty (CCP), a financial entity that is ideally very stable. One of the motivations for this is that trading with the CCP will mitigate the uncertainties that individual banks have about each other's qualities and thus help prevent counterparties from running during financial crisis. See [28] for an analysis of the clearing benefits of CCPs.

receive. However, a higher signal precision is harmful because it allows for a greater probability that the central agent becomes ostracized quickly, thus causing the network to fall apart. Such an event would prevent the generation of any future social welfare. Therefore there is a trade off between the initial mean and the signal precision of the central agent: it is desirable to have a central agent with a higher mean but a lower signal precision. In particular, choosing the agent based only on its initial mean expected quality is not optimal, whereas under complete information it would be optimal to always place the highest realized quality agent in the center.

We show these results formally in the next proposition. For concreteness, suppose that the central agent in the network is denoted by agent 1. The exogenous parameters of the agents are defined the same way as previously.

Proposition 7. *The overall social welfare is strictly increasing in* μ_1 *and strictly decreasing in* τ_1 *and* σ_1^2.

Figure 2 shows the trade-off between the mean and the signal precision of the central agent numerically via a simulation. It plots the contour lines of the overall *ex ante* welfare of the network, and it shows that social welfare increases as the initial mean increases and the signal precision decreases, and therefore selecting the best central agent depends on both factors[25].

We note that for the periphery agents on the other hand, the exogenous parameters have a much less clear relationship with the overall social welfare. We can show through examples that social welfare can increase or decrease in each of these factors for periphery agents. The same relationships as for the central agent could hold, and a simple example of this would be a two person network. However a marginally higher mean or a lower signal precision by a single periphery agent could also *decrease* overall welfare. For instance, consider a network where the central agent has an initial expected quality close to c, one

[25] The code used to generate the simulations in the paper can be found in the supplementary material.

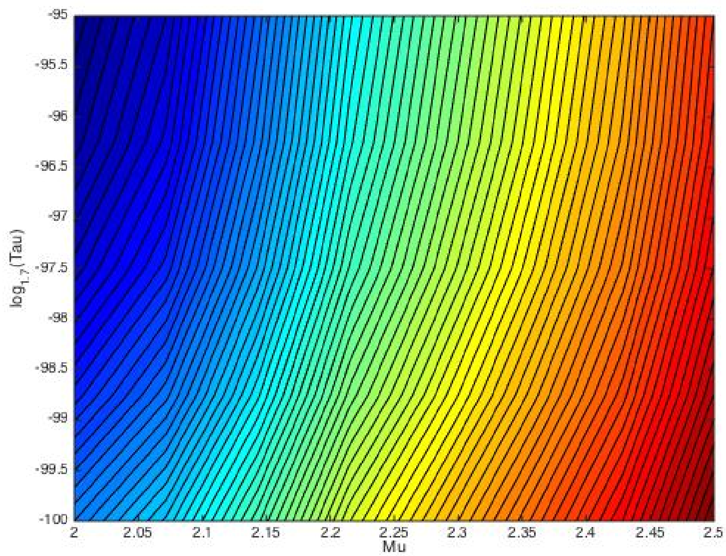

Figure 2: Simulation Illustrating Proposition 7

periphery agent denoted by agent i also has an initial expected quality close to c, and the qualities of all other periphery agents is very high. In such a case, increasing the expected quality of agent i by a small amount, or decreasing agent i's signal precision would harm overall social welfare. These changes would result in the central agent being connected to agent i for a longer stretch of time, which is undesirable since the other periphery agents are of much higher expected quality, and so causing the central agent to send more information is harmful. Therefore in such a network it would be better for agent i to send information more quickly in order for it to exit the network sooner.

We note that the trade-off identified above matters only if learning is fast enough, whereas if learning becomes very slow (or the designer becomes very impatient), then this trade-off goes away. This is summarized in the following proposition.

Proposition 8. *If the rate of learning becomes very slow (i.e. $\lambda \to 0$), then the optimal star network is obtained by placing the agent with the highest initial* *reputation in the center.*

Proposition 8 shows that the decision to place an agent at the center depends only on each agent's initial mean in the limit of very slow learning learning (a similar result holds for very high designer impatience). When the network is constrained to be a star network, the highest initial welfare is obtained by having the highest initial expected quality agent in the center if the learning is very slow.

8. Extensions

As seen above, learning can have a negative impact on social welfare in a variety of networks, and a large reason for this is the myopia of the agents. Since the agents are not experimenting for long enough, learning is inefficient and social welfare is lost. In this section, we consider four possible extensions that could alleviate this issue and allow for higher social welfare.

8.1. Linking Subsidy

A potential method of addressing the negative effects of learning is to give subsidies to the agents for linking with others. For instance, a company could give workers awards or bonuses for collaborating with colleagues. Or in a financial setting, a regulator may give financial incentives for firms conducting mutual investments, or guarantee interbank transactions during a financial crisis to lower default risk. We model a subsidy by assuming that for every link that an agent maintains, it receives an extra flow benefit of δ from the network designer. This linking subsidy does not affect the social welfare computation since it is a direct transfer from the network designer to the agent, but it would change agents' decisions of when to break a link. Since agents are myopic, an agent i will break its link with agent j if and only if agent j's reputation drops below $c - \delta$. The linking subsidy therefore causes the agents to learn more information about their neighbor's quality and break only if it is very likely to be bad. We show below that by properly choosing the linking subsidy the social welfare can improve compared with the case when there is no learning about agents' qualities. Let $W(\delta)$ denote the *ex ante* social welfare when the linking subsidy is equal to δ.

Theorem 5. *There exists $\bar{\bar{\delta}}$ such that $\forall \delta > \bar{\bar{\delta}}$, $W(\delta) > W^*$. Moreover, $\lim_{\delta \to \infty} W(\delta) = W^*$.*

Note that by Theorem 3, this result also shows that the social welfare is higher than the standard case with no subsidy. Thus by imparting subsidies on agents to encourage them to experiment for longer, the social welfare can be increased. The intuition is that when the link subsidy is high enough, any link that is broken will involve an agent that is of really low expected quality. Thus although the agent that is ostracized may still hurt from being disconnected, its neighbors will benefit by a sufficiently large amount that overall social welfare increases. Therefore learning is now beneficial and improves welfare overall. The second part of the theorem states that if the linking subsidy becomes too high, then the social welfare will converge to the social welfare without learning. This is because when the subsidy is too high it becomes almost impossible for a link to break, and so the network with high probability will not change, just like in the case without learning. Therefore having a linking subsidy is beneficial for the network, but the subsidy cannot be set too high either in order to maximize social welfare.

8.2. New Link Formation

Another way that learning would be more socially beneficial is if agents were able to form new links with other agents whose reputations are very high. In this extension, we assume that a pair of agents who are not initially linked according to the network constraint Ω can form a new link by incurring an instantaneous cost $\gamma > 0$. There is no cost to forming links with agents that they are connected to under Ω. So unlike previously when there was a hard barrier between agents not connected according to Ω, agents can now break this barrier by paying an instantaneous cost. This cost could be exogenous, for instance the cost of investing time and energy into becoming familiar with a new agent, or the cost of reducing some physical barrier between the agents (distance or geographic barriers). The cost could also be set by the network designer such as a tax on link creation. Since we assume the formation cost is instantaneous, it is infinitesimal in the social welfare calculation and so only affects welfare through its impact on agent actions.

We assume that forming a link this way requires bilateral consent as usual. Agent i will want to form a link with agent j if agent j's reputation is higher than $c+\gamma$. Therefore a new link between agents i and j is formed at time t if and only if $\mu_i^t \geq c+\gamma$ and $\mu_j^t \geq c+\gamma$. The dynamics of our model will now feature some agents attaining high reputation levels and being able to link with other previously inaccessible agents that have also attained high reputation levels. Allowing these two high expected quality agents to link together will improve social welfare due to the large mutual benefits that are generated from their link.

We can compare the social welfare produced by allowing this extra link formation against the social welfare in the basic model. Let $W(\gamma)$ denote the *ex ante* social welfare when the link formation cost is equal to γ, and let W be the social welfare in the basic model without the extra link formation.

Theorem 6. *There exists $\bar{\gamma}$ such that $\forall \gamma \geq \bar{\gamma}$, $W(\gamma) > W$.*

This theorem states that if the link formation cost is high enough then the social welfare is improved over the base model because two agents that decide to form a new link will do so with very high reputations. Thus the social welfare generated by a new link is likely to be very high as well, and this dominates any potential informational externalities that the link could create. Note however that a γ that is too low may actually harm welfare. For instance suppose according to Ω there are a group of moderate expected quality agents that are all linked to a very high expected quality agent but separated from each other. This is similar to the core-periphery setting examined in Theorem 4. In such a case, allowing moderate reputation agents to link with each other would cause them to harm each other via the negative informational effects of the link. This would reduce welfare overall compared to the base model. Therefore allowing for new link formation can improve welfare, but the threshold for the link being formed must be sufficiently high as well. The optimal $\bar{\gamma}$ would depend on the specific properties of the network. If as in the example there exists a group of very high reputation agents that the moderate reputation agents are linked with, then $\bar{\gamma}$ would likely be higher as well, as it becomes more important for moderate reputation agents to not be linked with each other.

8.3. Agent Entry

Our model can also be tractably extended to allow agents to enter into the network over time. Specifically, suppose that for the set of N agents in V there is a corresponding set of entry times $\{e_i\}_{i \in V}$, with $e_i \geq 0 \; \forall i$. Agents with $e_i = 0$ are present in the network at the beginning, while agents with $e_i > 0$ enter later on. These entry times are fixed and known to the agents in the model. The network constraint Ω is now defined over the set of all N potential agents and still specifies which agents are allowed to connect to each other, including agents that arrive later. This network constraint effectively determines where agents enter into the network at their entry times. The learning process is the same as before, with learning occurring for agents within the network based on their current amount of neighbors, and no learning occurring for an agent that has not yet entered.

Agents still make decisions myopically and will connect with a neighbor for as long as that neighbor's reputation is above the connection cost. Since we assume all agents have initial reputations above the cost, an incumbent agent will always wish to connect with a newly entering agent. However, the new agent would not want to connect with one of its neighbors that has already been ostracized previously within the network. The dynamics will evolve similarly to before, with agents connecting to neighbors until a neighbor's reputation falls too low, at which point the neighbor will be ostracized. The difference now is that new agents will arrive at certain times, and when they do they will change the benefits and amount of information produced by the network.

We can compare the model with agent entry against the base model where all agents were present in the beginning, i.e. $e_i = 0 \; \forall i \in V$. We fix a network constraint Ω and perform comparative statics on the entry times of the agents. We first show that incorporating agent entry will not change either the set or the distribution of stable networks.

Proposition 9. *The set of stable networks is unchanged with agent entry. The probability of each stable network emerging is the same as that given in Corollary 2 and identical to the case without agent entry.*

Although the properties of the final stable networks are not affected by agent entry, the overall social welfare will be affected. It is possible to calculate social welfare in a similar method as in Theorem 2, as we can account for agent entry by rescaling the hitting times of the agents in the network appropriately. Incorporating agent entry has two separate effects on social welfare: first, the links that the entering agent has are started later, so the benefits from those links are realized later as well and thus discounted more heavily; second, the neighbors of the entering agent send less information before that agent enters, and the agent itself may send information more slowly if one of its neighbors is ostracized before it enters, delaying the time at which the agent and its neighbors are potentially ostracized from the network. The first effect hurts social welfare because the benefits from any link are positive in expectation. However, the second effect can improve social welfare by delaying the agents' ostracization times and increasing the benefits that each agent is able to extract from the network. It is possible for the second effect to dominate the first, so that delaying entry for an agent raises social welfare overall.

Theorem 7. *For some network parameters, increasing a single agent's entry time e_i can increase social welfare.*

Figure 3 shows the trade-off explicitly via a simulation that highlights Theorem 7. When a new agent of moderate reputation enters later, social welfare

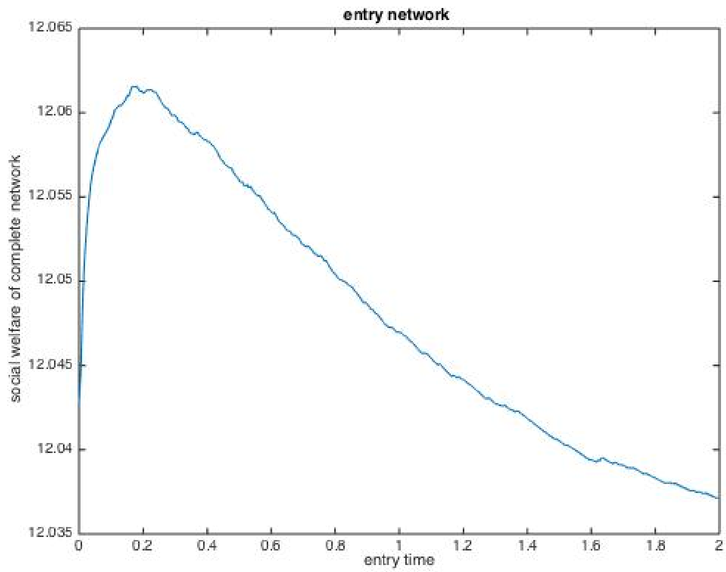

Figure 3: Simulation Illustrating Theorem 7

initially increases because the incumbent agents have more time to benefit from their links. However, if the entry time becomes too large then the social welfare decreases, since the reputations of the incumbent agents have stabilized already, and it is thus better to have the new agent enter sooner and benefit from the network as well.

As an implication, a financial regulator may wish to delay new firms from entering the network in times of crisis when there is a lot of uncertainty, and then allow them to enter once the crisis has ended and reputations are more stable. Or in an organization, a firm may wish to not expand too quickly, and instead take the time to allow the current workers to better understand each other first.

9. Conclusion

This paper analyzed agent learning and the resulting network dynamics when there is incomplete information. We presented a highly tractable model that explicitly characterized what the set of stable networks are for a given network, showed how learning affects both individual and social welfare depending on the specific network topology, and analyzed what optimal initial network structures look like for different groups of agents. Our results shed new light on network dynamics in real world situations, and they offer guidelines for optimal network design when there is initial uncertainty about the agents. When agents are sufficiently myopic in their actions, ostracism becomes harmful not just for the ostracized agents themselves, but to all agents in an *ex ante* fashion. A network designer should thus structure links appropriately in order to minimize the negative effects of ostracism.

Our results could be extended in several interesting ways. One natural extension would be to allow the qualities of agents to evolve over time. In the simplest extension, the agent's true quality q_i itself change according to an exogenous stochastic process, for instance a Brownian motion. More interestingly, it would be natural to assume that the evolution of true quality depends endogenously on the information the agent receives so that agents who receive better information tend to develop higher true qualities and hence also generate better information in the future. Thus, the structure of the network and the true qualities of the agents in the network co-evolve. Higher reputation agents may link to agents that are also of higher reputation, and so their true qualities would improve as well, while lower reputation agents may struggle to find good agents to link with, and their true qualities would decline as a result.

It would also be interesting to allow agents to engage in games with their linked neighbors instead of merely generating flow benefits. Games played over networks have been analyzed in several papers within the networks literature (see [29] for a review), but never in a dynamic setting with learning such as that considered in the current paper. The game played by agents could be a prisoner's dilemma or another type of cooperation game where the payoffs depend on the agent's types. Agents would need to seek out other agents that they can achieve high payoffs in the game with, and this process would also require learning over time about a neighbor's type. As agents are able to learn each other's type more accurately, they may achieve greater efficiency in their plays and also sustain cooperation for a longer length of time.

Appendix

Proof of Proposition 1

Proof. Suppose for now that agent i's reputation always evolves at the constant signal precision τ_i. Then given the true quality q_i for agent i, the probability that agent i's reputation never hits c before t can be found using standard arguments (see for example [30]) and is given by

$$P(S_i^t|q_i) = \Phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right)$$
$$- \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c))\Phi\left(\sqrt{t\tau_i}(q_i - c) - \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right)$$

Therefore, given q_i, the probability that agent i stays in the network is

$$P(S_i|q_i) = \lim_{t\to\infty} P(S_i^t|q_i) \tag{.1}$$

- If $q_i > c$, as $t \to \infty$, then we have $\Phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \to 1$ and $\Phi\left(\sqrt{t\tau_i}(q_i - c) - \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \to 1$. Thus, $P(S_i|q_i) = 1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c))$, namely agent i stays in the network with positive probability and the probability is increasing in the true quality q_i.

- If $q_i < c$, as $t \to \infty$, then we have $\Phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \to 0$ and $\Phi\left(\sqrt{t\tau_i}(q_i - c) - \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \to 0$, thus $P(S_i|q_i) = 0$, namely agent i's reputation hits c before $t = \infty$ for sure.

- If $q_i = c$, it is clear that $P(S_i^t|q_i) = 0$ as $t \to \infty$.

Taking the expectation over q_i, we have

$$P(S_i) = \int_c^\infty (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i/\sigma_i \tag{.2}$$

From the above expression we can see that $P(S_i)$ only depends on the initial quality distribution (μ_i and σ_i) and the link cost c but does not depend on the Brownian motion precision τ_i. Since breaking links only changes the Brownian motion precision, the probability that an agent's reputation never hits c is independent of the initial network G^0 or the signal precision τ_i. $\square$

Proof of Corollary 1

Proof. We first show that $P(S_i)$ is increasing in μ_i. Let $q_i - \mu_i = x$. Then $P(S_i)$ can be rewritten as

$$P(S_i) = \int_{c-\mu_i}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(\mu_i - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx/\sigma_i$$

Consider a larger expected quality $\mu_i' > \mu_i$, we have

$$\begin{aligned} P(S_i|\mu_i') &= \int_{c-\mu_i'}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i' - c)(\mu_i' - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx/\sigma_i \\ &> \int_{c-\mu_i}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i' - c)(\mu_i' - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx/\sigma_i \\ &> \int_{c-\mu_i}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(\mu_i - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx/\sigma_i = P(S_i|\mu_i) \end{aligned}$$

Therefore, $P(S_i)$ is increasing in μ_i.

Next we show that $P(S_i)$ is decreasing in σ_i.

$$\begin{aligned} P(S_i) &= \int_{c-\mu_i}^{\infty} \phi(\frac{x}{\sigma_i})dx/\sigma_i - \int_{c-\mu_i}^{\infty} e^{-\frac{2}{\sigma_i^2}(\mu_i - c)(\mu_i - c + x))} \frac{1}{\sigma_i\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x}{\sigma_i}\right)^2} dx \\ &= \int_{c-\mu_i}^{\infty} \phi(\frac{x}{\sigma_i})dx/\sigma_i - \int_{c-\mu_i}^{\infty} \frac{1}{\sigma_i\sqrt{2\pi}} e^{-\frac{1}{2\sigma_i^2}(2(\mu_i - c) + x)^2)} dx \\ &= \int_{c-\mu_i}^{\infty} \phi(\frac{x}{\sigma_i})dx/\sigma_i - \int_{\mu_i - c}^{\infty} \phi(\frac{x}{\sigma_i})dx/\sigma_i = \int_{c-\mu_i}^{\mu_i - c} \phi(\frac{x}{\sigma_i})dx/\sigma_i \end{aligned}$$

Therefore, $P(S_i)$ is decreasing in σ_i.

Finally, we show that $P(S_i)$ is decreasing in c. Consider a smaller $c' < c$, we have

$$\begin{aligned} P(S_i|c) &= \int_{c}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i/\sigma_i \\ &< \int_{c}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c')(q_i - c')))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i/\sigma_i \\ &< \int_{c'}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c')(q_i - c')))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i/\sigma_i = P(S_i|c') \end{aligned}$$

The first inequality is because for $q_i > c$, $1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)) < 1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c))$. The second inequality is because for $c' < q_i < c$, $1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)) > 0$. $\square$

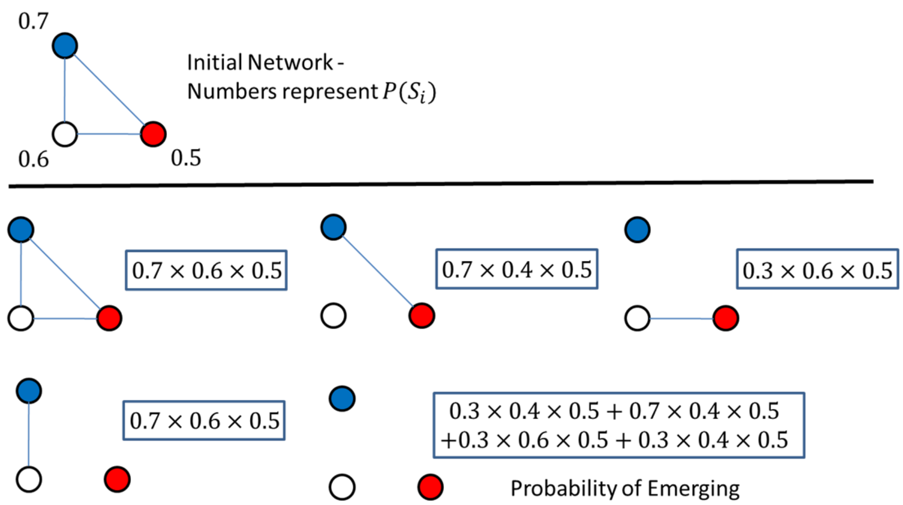

Figure A.1: Set of Stable Networks Given an Initial Network: This figure shows the five possible stable networks that could emerge given an initial network of three agents. In addition, $P(S_i)$ is given for all the agents, which allows us to calculate the exact probability of each of these networks emerging. For the first four networks, there is only one realization of $\{S_i, ^\neg S_i\}_{i \in V}$ that corresponds to it. For the last network, there are four possible realizations, one in which $^\neg S_i$ occurs for all agents, and three in which S_i occurs for a single agent.

Proof of Theorem 1

Proof. It is clear that a network G must be a subset of G^0 and can be stable if and only if there exists at least one combination of events $\{S_i, ^\neg S_i\}_{i \in V}$ such that $g_{ij} = I\{S_i\}I\{S_j\}I\{g^0_{ij} = 1\}$. Thus the set of stable networks does not depend on the learning speed. Moreover, according to Proposition 1, $P(S_i)$ is independent over the different agents and does not depend on the speed of learning. Hence the probability that any specific link exists in the stable network exists also independent of the learning speed, so the probability of any stable network emerging is also independent of the learning speed. □

Proof of Lemma 1

Proof. Since the Brownian motion precision is constant, using the survival probability

$$P(S_i^t|q_i) = \Phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \\ - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c))\Phi\left(\sqrt{t\tau_i}(q_i - c) - \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right)$$

we can compute $f(\hat{\varepsilon}_i^t|q_i) = -\frac{dP(S_i^t|q_i)}{dt}$ as

$$f(\hat{\varepsilon}_i^t|q_i) = -\frac{1}{2}\left(\sqrt{\tau_i}(q_i - c)t^{-1/2} - \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \\ + e^{-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)}\frac{1}{2}\left(\sqrt{\tau_i}(q_i - c)t^{-1/2} + \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i - c) - \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \\ = -\frac{1}{2}\left(\sqrt{\tau_i}(q_i - c)t^{-1/2} - \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \\ + \frac{1}{2}\left(\sqrt{\tau_i}(q_i - c)t^{-1/2} + \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \\ = \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right)$$

Taking the expectation over q_i, we obtain $f(\hat{\varepsilon}_i^{t_i})$. □

Algorithm 1 Hitting Time Mapping Function

Input: $\boldsymbol{t}^B$, base precision $\tau_i, \forall i$ and initial graph G^0.

Output: new hitting time vector $\boldsymbol{t}$.

Initiate: $d_i = t_i^B \tau_i$, $v_i = k_i^0 \tau_i$

Initiate: $\mathcal{N} = \{i : t_i^B < \infty\}$, $t_i^{G^0} = 0, \forall i \in \mathcal{N}$ and $t_i^{G^0} = \infty, \forall i \notin \mathcal{N}/$

while $\mathcal{N} \neq \emptyset$ **do**

  Let $i^* = \min_{i \in \mathcal{N}} d_i / v_i$.

  Update $t_i^{G^0} := t_i^{G^0} + d_{i^*}/v_{i^*}, \forall i \in \mathcal{N}$.

  Update $d_i := d_i - v_i \times d_{i^*}/v_{i^*}$.

  Update $\mathcal{N} := \mathcal{N}/i^*$.

  Update $k_i = \max\{1, k_i - 1\}$, for all i such that $G^0_{ii^*} = 1$.

  Update $v_i = k_i \tau_i$

end while

Figure A.2: Algorithm for Computing Hitting Time Mapping Function M

Proof of Theorem 3

Proof. Consider the *ex ante* surplus W_{ij} that agent i obtains from the link with a neighbor j. The *ex ante* welfare for agent i is simply the summation of this surplus over all j that i is linked with. W_{ij} can be computed as

$$W_{ij} = \int_q \int_0^\infty e^{-\rho t} P(L_{ij}^t | q)(q_j - c) dt \phi(q) dq$$

where $P(L_{ij}^t|q)$ is the probability that the link between i and j still exists at time t. Let t^* be the time at which the link between i and j is broken. Then the social welfare can be computed as

$$\begin{aligned} W_{ij} &= \int_q \int_0^\infty e^{-\rho t}(q_j - c) dt \phi(q) dq - E_{t^*}[\int_{t^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t^*) dt] \\ &= \int_0^\infty e^{-\rho t}(\mu_j - c) dt - E_{t^*}[\int_{t^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t^*) dt] \end{aligned}$$

where the expectation is taken over the realizations in which the hitting time is t^*. The second term can be further decomposed. Let t_i^* denote the case when

$t^* = t_i$, namely agent i's reputation hits c before agent j, and t_j^* be the case where $t^* = t_j$, namely agent j's reputation hits c before agent i. Then

$$\begin{aligned} W_{ij} = & \int_0^\infty e^{-\rho t}(\mu_j - c)dt \\ & - E_{t_i^*}[\int_{t_i^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t_i^*)dt] - E_{t_j^*}[\int_{t_j^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t_j^*)dt] \end{aligned}$$

In the case of t_j^*, for any $t \geq t_j^*$, since the learning has stopped, $E_{q_j}(q_c - c | t \geq t_j^*) = 0$ by the definition of t_j^*. Similarly, $E_{q_j}(q_j - c | t \geq t_i^*) > 0$ because at t_i^* the expected quality of q_j is strictly greater than c since agent j has not been ostracized. Therefore,

$$W_{ij} = W_{ij}^* - E_{t_i^*}[\int_{t_i^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t_i^*)dt] < W_{ij}^*$$

Summing over all j that i is linked with, we conclude that the agent i's *ex ante* welfare with learning is strictly less than that when there is no learning. Note that this result holds independently of the values of $\tau_1, ..., \tau_N$, as the signal precisions affect only the distribution of agent hitting times, but not the expected quality of the agents conditional on ever being ostracized. □

Proof of Proposition 2

Proof. Consider any *ex post* realization $\varepsilon = \{\varepsilon_i^{t_i}\}_{i \in V}$. If $t_i = \infty$, then changing τ_i alone does not change the fact that agent i would stay in the network forever, as so it does not affect the hitting time realization of any other agent either. Therefore agent i's welfare $W_i(\varepsilon)$ is not affected. If $t_i < \infty$, then the welfare of agent i depends on (1) the expected quality of all the neighboring agents j whose $t_j = \infty$ and (2) its own hitting time t_i. Since (1) is not affected by changing τ_i, we only need to study how τ_i affects t_i.

Intuitively t_i is decreasing in τ_i since agent i's information sending speed is faster due to a higher precision. We provide a more rigorous proof by contradiction as follows. Suppose agent i's new hitting time increases to $t_i' = t_i + \Delta > t_i$. In this new realization, consider the duration from 0 to t_i. Since $t_i' > t_i$, all other agents' information sending process and speed do not change before t_i. Hence, agent i's instantaneous precision at $t \leq t_i$ changes to $(\tau_i^t)' = \frac{\tau_i'}{\tau_i}\tau_i^t$. Hence, information sending by agent i is faster at any moment in time before t_i. Since, the stopping time t_i' is larger than t_i, the total amount of information sent by agent i given τ_i' is larger than that given τ_i. Because the total information sent should remain the same, this causes a contradiction. Therefore t_i' should be smaller than t_i for a larger τ_i'. □

Proof of Proposition 3

Proof. Consider any *ex post* realization $\varepsilon = \{\varepsilon_i^{t_i}\}_{i \in V}$. If $t_i = \infty$, then increasing agent i's signal precision τ_i does not change the realization $\varepsilon_i^{t_i}$. Hence t_j is not affected. If $t_i < \infty$, then according to Proposition 1, the new hitting time t_i' is sooner if agent i's signal precision is larger. This causes the link between agent i and j to be severed (weakly) sooner, leading to a (weakly) later hitting time of agent j because agent j will send information at a slower speed for a longer time. Since changing agent i's signal precision does not change the finiteness of the hitting time of all other agents, agent j's welfare increases due to a longer hitting time for itself. □

Proof of Proposition 4

Proof. We prove by constructing a counterexample, which is shown in Figure A.3. Consider a network with $K > 3$ agents. Agents 1, 2, 3 form a line and the other $K - 3$ agents connect to both and only agents 1 and 2. We assume that agent 3's true quality is perfectly known (initial variance 0) and large. Hence, agent 3's reputation never hits c. We also assume that the mean qualities of agents 4 to K are close to c. Hence, agent 2 almost does not gain benefit from those agents even when $K \to \infty$.

Consider a realization in which agent 1's reputation hits c at $t_1 < \infty$ and agent 2's reputation hits c at $t_2 < \infty$. By increasing the signal precision of agent 1, its hitting time decreases to $t_1' < t_1$. If $t_1' > t_2$, then agent 2's hitting time is not affected, i.e. $t_2' = t_2$. Otherwise, the new hitting time may be different from t_2. To simplify the analysis, we consider the extreme case in

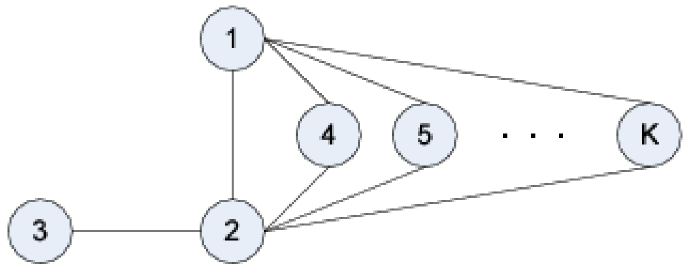

Figure A.3: Counterexample for Proposition 4

which $\tau_i \to \infty$, thereby $t_1' \to 0$. Therefore, agent 2 loses the link with agent 1 from the beginning in any realization. However, since agents 3 to K also lose the link with agent 1 from the beginning, for those whose hitting time was earlier than t_2, their hitting time would increase by a factor of 2. If there are at least three agents among 4 to K whose hitting was between $[t_2/4, t_2/2]$, agent 2's information sending speed will increase sufficiently much that agent 2's hitting time is smaller. By making K large we can always making the probability of this event be large enough. Thus, agent 2's hitting time will decrease on average. □

Proof of Proposition 5

Proof. From Theorem 2, we have that the *ex ante* social welfare is given by:

$$W = E_{\hat{\varepsilon}} \sum_i \left(\frac{1 - e^{-\rho M_i(t)}}{\rho} \sum_{j: g_{ij}^0 = 1, t_j = \infty} \frac{\mu_j - c}{P(S_j)} \right)$$

If the designer is completely impatient, it only cares about the social surplus at time 0 since with probability approaching 1 no links will be broken among the agents. Since all agents' expected qualities are above the linking cost, having all agents connected with each other yields the highest social surplus. Similarly if learning becomes very slow, then the agent's reputations are never updated and the same reasoning applies. In both cases the $e^{-\rho M_i(t)}$ term approaches zero in the above equation regardless of the network structure, and so adding more agents increases welfare.

If the designer is completely patient, only the stable networks matter. Since the stable network does not depend on the speed of learning and the probability that an agent stays in the stable network is independent of others by Proposition 1, having all agents connected with each other leads to the maximum number of links in the stable networks and hence the highest social surplus. Similarly if learning becomes very fast, the stable network will always be reached immediately and the same reasoning applies. In both cases the $e^{-\rho M_i(t)}$ term approaches one in the above equation regardless of the network structure, and so adding more agents increases welfare. ☐

Proof of Proposition 6

Proof. We will prove that for $\bar{\mu}$ large enough, the social welfare of any non fully connected network will be increased through the addition of any new link. Therefore the welfare of the fully connected network will be greater than the welfare of any other network. Consider an arbitrary network constraint Ω that is not fully connected. Suppose that a link between agents i and j is added to the network, and consider the welfare under the new network constraint Ω'.

First consider the change in welfare of agent i. In any realization where agent i is ostracized, its welfare through having the extra link with j decreases by no more than $\frac{(N-2)\mu}{\rho}$, the welfare loss when it loses all its links with the other agents immediately. In any realization where agent i is not ostracized, its welfare with the additional link increases by $\frac{\mu}{\rho}$, the discounted value of the new link given the expected quality of agent j. Thus the change in welfare for agent i is bounded below by $P(S_i)\frac{\mu}{\rho} + (1 - P(S_i))\frac{(N-2)\mu}{\rho} = \mu(P(S_i)\frac{(N-1)}{\rho} - \frac{(N-2)}{\rho})$. Similarly, we can show that the change in welfare for agent j is bounded below by $\mu(P(S_j)\frac{(N-1)}{\rho} - \frac{(N-2)}{\rho})$.

Now consider the change in welfare for all the other agents in the network. In any realization where both agent i and agent j are not ostracized, the hitting times of all the agents in the network are unaffected by the new link. In any realization where either agent i or agent j are ostracized, the change in welfare for all the other agents is bounded below by $\frac{(N-2)(N-1)\mu}{\rho}$. Thus the total change in welfare for all other agents in the network is bounded below by $[P(S_i)(1 - P(S_j)) + (P(S_j)(1 - P(S_i)) + (1 - P(S_i))(1 - P(S_j))]\frac{(N-2)(N-1)\mu}{\rho}$.

Combining the above two observations, we note that the change in welfare for the whole network is bounded below by $\mu[P(S_i)\frac{(N-1)}{\rho} - \frac{(N-2)}{\rho} + P(S_j)\frac{(N-1)}{\rho} - \frac{(N-2)}{\rho} + P(S_i)(1-P(S_j)) + (P(S_j)(1-P(S_i)) + (1-P(S_i))(1-P(S_j))\frac{(N-2)(N-1)}{\rho}]$. When $\bar{\mu}$ is large, $P(S_i)$ converges to 1 by Proposition 1. Thus for $\bar{\mu}$ large enough, the lower bound for the change in welfare of agents i and j converges to $\frac{2(N-1)\mu}{\rho}$, a positive number.

When $\bar{\mu}$ is large, $P(S_i)$ and $P(S_j)$ converge to 1 by Proposition 1. Therefore the lower bound for the change in welfare converges to $\frac{2\mu}{\rho}$, a positive. □

Proof of Theorem 4

Proof. We first show that all high type agents should connect to all other high type agents. This is based on a similar argument as in the proof of Proposition 6. Since when $\mu_H \to \infty$, all high type agents will stay in the stable network with very high probability, adding a link between any two high type agents will strictly improve their welfare while impacting the welfare of all other agents with very low probability. Hence, there must exist a large enough value for μ_H such that the welfare of high type agents is maximized when all high type agents connect to all other high type agents in the initial network.

Next we show that all low type agents should not connect to each other in any network where each is linked to at least 1 high type agent. When $\mu_H \to \infty$, the welfare obtained by a link with any low type agent j is dominated by that a link with high type agents, i.e. we can suppose that the welfare received by a link with another low type agent is approximately zero in comparison to a link with the high type agents. Having additional links with other low-type agents reduces the hitting time of agent j, $M_j(t)$, in the event that it gets ostracized, thereby reducing agent j's welfare by more than the welfare gain of the additional link. Therefore, low type agents do not connect to each other in the optimal initial network.

Finally we show that all low type agents should connect with every high type agent. Since the probability that the high type agent is ostracized approaches zero, such a link does not affect them relative to the extra welfare that the low type agents receive. Therefore we consider only the effect on the welfare of the low type agent to be connected with all high type agents. In a realization where the low type agent is not ostracized, this is optimal for all agents, as the high type agent stays in the network with very high probability when μ_H is large enough. Thus both agents have their welfare increased while not affecting the welfare of all other agents. We show that it is also optimal in realizations where the low type agent is ostracized. Again we will assume that the high type agent is not ostracized, which will hold for μ_H high enough. The low type agent receives a flow payoff of μ_H from every high type agent that it has an active link with. Note that in the hitting time mapping function the hitting time of an ostracized agent i is scaled by $1/K$, where K is the total number of high type neighbors. Thus the decrease in hitting time is exactly balanced out by the increase in flow payoff in the case without discounting, and with discounting it is strictly better for the low type agent to have an extra link. □

Proof of Proposition 7

Proof. We can break social welfare into two components: the welfare of the central agent, and the welfares of each periphery agent. Notice that the welfares of the periphery agents are strictly increasing in μ_1 but do not depend on σ_1 or τ_1 for similar reasons as in the proof of Theorem 3. Also, the welfare of the central agent is strictly increasing in μ_1 as that allows the central agent to stay in the network for a longer period of time. Thus overall social welfare is increasing in μ_1. The welfare of the central agent is strictly decreasing in τ_1 for the same reasons as in Proposition 2. Thus overall social welfare is decreasing in this parameter. □

Proof of Proposition 8

Proof. In the limit of very slow learning, only the initial welfare generated matters, and placing the agent with the highest expected quality in the center generates the highest welfare over all initial network structures. □

Proof of Theorem 5

Proof. (1) Consider the welfare on link l_{ij}. As in the proof of Theorem 3, let t_i^* denote the event in which agent i's reputation hits $c - \delta$ at time t_i^* before agent j. The *ex ante* welfare of link l_{ij} can be computed as

$$W_{ij} + W_{ji} = \int_0^\infty e^{-\rho t}(\mu_i^0 + \mu_j^0 - 2c)dt$$
$$-E_{t_i^*}[\int_{t^*}^\infty e^{-\rho t} E_{q_i,q_j}(q_i + q_j - 2c|t \geq t_i^*)dt]$$
$$-E_{t_j^*}[\int_{t^*}^\infty e^{-\rho t} E_{q_i,q_j}(q_i + q_j - 2c|t \geq t_j^*)dt]$$

Note that the first integral in the above equation represents $W_{ij}^* + W_{ji}^*$, the social welfare of the link without learning.

In the case of t_i^*, for any $t \geq t_i^*$, since the learning has stopped, $E_{q_i,q_j}(q_i - c|t \geq t_i^*) = c - \delta - c = -\delta$ by the definition of t^*. Since agent j is not ostracized, we would have $E_{q_i,q_j}(q_j - c|t \geq t_i^*) > -\delta$. Let $h(\delta, t_i^*) = E_{q_i,q_j}(q_i + q_j - 2c|t \geq t_i^*) = E_{q_i,q_j}(q_j|t \geq t_i^*) - 2c - \delta$. This is the net change in flow payoff after the link is severed. We will show that for any t_i^*, $h(\delta, t_i^*) < 0$ if δ is sufficiently large. A symmetric argument then establishes that $h(\delta, t_j^*) < 0$, and the two together imply that the welfare of the link with learning is greater than $W_{ij}^* + W_{ji}^*$. Then, adding up over all links shows that the overall social welfare is higher than that without learning.

To prove that $h(\delta, t_i^*) < 0$ if δ is sufficiently large, we will show that $E_{q_i,q_j}(q_j|t \geq t_i^*)$ is bounded above for any t_i^* as δ tends to infinity. Consider any *ex post* realization of t_i^*, which implies that agent j's reputation does not hit $c - \delta$ before t_i^*. There are two possibilities for agent j's reputation (here we assume that agent j continues sending information at its fixed signal precision if all its other neighbors are ostracized, as in section 4):

- ζ_1: it never hits $c - \delta$ after t_i^* either.
- ζ_2: it hits $c - \delta$ at some time after t_i^*.

Clearly, $E(q_j|\zeta_1) > E(q_j|\zeta_2) = c - \delta$. Hence $E_{q_i,q_j}(q_j|t \geq t^*) < E(q_j|\zeta_1)$. The value of $E(q_j|\zeta_1)$ is given by equation 5 in the text, with c replaced by $c - \delta$.

When $\delta \to \infty$, using equation 5 we can show that $\lim_{\delta\to\infty} E(q_j|\zeta_1) = \mu_j^0$ through the application of L'Hopital's rule.

Therefore, $\forall \epsilon > 0$, there exists δ'_{ij} such that $\forall \delta > \delta'_{ij}$, $E(q_j|\zeta_1) - \mu_j^0 < \epsilon$. Hence, fix a value of $\epsilon > 0$ and let $\bar{\delta}_{ij} = \max\{\delta'_{ij}, \mu_j^0 - 2c + \epsilon\}$, which ensures for all $\delta > \bar{\delta}_{ij}$, $E(q_j|\zeta_1) - 2c - \delta < 0$. This also implies that $h(\delta, t_i^*) < 0$ for all t_i^* and $\delta > \bar{\delta}_{ij}$. By choosing $\bar{\delta} = \max_{i,j} \bar{\delta}_{ij}$, we ensure the overall *ex ante* social welfare is greater than W^*.

(2) Define $H_{ij}(\delta) = E_{t_i^*}[\int_{t_i^*}^{\infty} e^{-\rho t} E_{q_i,q_j}(q_i + q_j - 2c|t \geq t_i^*)dt]$. We will prove $\lim_{\delta\to\infty} H_{ij}(\delta) = 0$. To prove this, we will show that for any sequence $\delta_n \to \infty$, the sequence $H_{ij}(\delta_n) \to 0$. We divide $H_{ij}(\delta)$ into two parts,

$$\begin{aligned} H_{ij}(\delta) = E_{t_i^* < \hat{t}(\delta)} &\left[\int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c|t \geq t_i^*)]dt\right] \\ + E_{t_i^* \geq \hat{t}(\delta)} &\left[\int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c|t \geq t_i^*)]dt\right] \\ &= H'_{ij}(\delta) + H''_{ij}(\delta) \end{aligned}$$

for some $\hat{t}(\delta)$. We will find a sequence $\hat{t}(\delta_n)$ such that both $H'_{ij}(\delta_n) \to 0$ and $H''_{ij}(\delta_n) \to 0$ as $\delta_n \to \infty$.

Let $\hat{t}(\delta_n) = \delta_n$. First we will show that for δ_n large enough, $P(t_i^* < \delta_n) < \frac{1}{\delta_n^2}$. Note for a given q_i, the probability that the agent is ostracized before time δ_n is equal to:

$$\begin{aligned} 1 - P(S_i^{\delta_n}|q_i) = 1 - \Phi&\left(\sqrt{\delta_n \tau_i}(q_i - c + \delta_n) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c + \delta_n)}{\sqrt{\delta_n \tau_i}}\right) \\ -\exp(-\frac{2}{\sigma_i^2}(\mu_i - c + \delta_n)(q_i - c + \delta_n))\Phi&\left(\sqrt{\delta_n \tau_i}(q_i - c + \delta_n) - \frac{\frac{1}{\sigma_i^2}(\mu_i - c + \delta_n)}{\sqrt{\delta_n \tau_i}}\right) \end{aligned}$$

Note that $\lim_{x\to\infty} \Phi(x) = 1 - \frac{e^{\frac{-x^2}{2}}}{x\sqrt{2\pi}}$. Therefore the term above approaches zero faster than $\frac{1}{\delta_n^2}$ as $\delta_n \to \infty$. Integrating over all q_i shows that $P(t_i^* < \delta_n) < \frac{1}{\delta_n^2}$ for large δ_n.

Now consider $H'(\delta_n)$, it is bounded by

$$
\begin{aligned}
|H'(\delta_n)| &< P(t_i^* < \delta_n) \sup_{t_i^* < \delta_n} | \int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c|t \geq t_i^*)]dt| \\
&< \frac{\sup_{t_i^* < \delta_n} | \int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c|t \geq t_i^*)]dt|}{\delta_n^2} \\
&< \frac{1}{\rho \delta_n^2} \sup_{t_i^* < \delta_n} |E[q_j|\zeta_1] - c + \delta_n|
\end{aligned}
$$

Since as $\delta_n \to \infty, E[q_j|\zeta_1] \to \mu_j^0$, we conclude that $|H'(\delta_n)| \to 0$.

Consider $H''(\delta_n)$, it is bounded by

$$
\begin{aligned}
|H''(\delta_n)| &< \sup_{t_i^* \geq \delta_n} | \int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c|t \geq t_i^*)]dt| \\
&< \frac{1}{e^{\rho \delta_n}} \sup_{t_i^* > \delta_n} | \int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho (t-\delta_n)}(q_i + q_j - 2c|t \geq t_i^*)]| \\
&< \frac{1}{\rho e^{\rho \delta_n}} \sup_{t_i^* < \delta_n} |E[q_j|\zeta_1] - c + \delta_n|t \geq t_i^*|
\end{aligned}
$$

Similarly, since as $\delta_n \to \infty, E[q_j|t \geq t_i^*] \to \mu_j^0$, we conclude that $|H''(\delta_n)| \to 0$. □

Proof of Theorem 6

Proof. Consider any realization ε when link formation is not allowed. The *ex post* welfare $W(\varepsilon)$ is changed only when there is some time t^* such that there exist two agents i and j, who are not initially connected, such that $\mu_i^{t^*} \geq c + \gamma$ and $\mu_j^{t^*} \geq c + \gamma$. In the original realization ε, conditional on t^*, there are two cases

- ζ_1: Both agents' reputations never hit c after t^*.
- ζ_2: At least one agent's reputation hits c after t^*.

When ζ_2 occurs, allowing link formation may change the hitting time of all agents' in the network and hence, the welfare $W(\varepsilon|\zeta_2)$ may change. However, the probability of ζ_2 occurring tends to zero as $\bar{\gamma}$ tends to infinity by Proposition 1. When ζ_1 occurs, the social welfare increases by at least $\frac{e^{-\rho t^*}}{\rho} \frac{(c+\bar{\gamma})-(1-P(\zeta_1))c}{P(\zeta_1)}$. When ζ_2 occurs, the welfare decreases by at most $B(\zeta_2)$, a function that is at most linear in $\bar{\gamma}$ as it grows large, since the set of agents and their initial qualities are fixed. Thus the overall change in welfare can be written as

$$W'(\varepsilon) - W(\varepsilon) \geq P(\zeta_1)\frac{e^{-\rho t^*}}{\rho}\frac{(c+\bar{\gamma}) - (1-P(\zeta_1))c}{P(\zeta_1)} - P(\zeta_2)e^{-\rho t^*}B(\zeta_2)$$

By choosing $\bar{\gamma}$ large enough, we can ensure that $P(\zeta_2)$ is small enough such that the change is positive in all such realizations ε. Therefore $W(\gamma) > W$.

□

Proof of Proposition 9

Proof. First note that Proposition 1 still holds for each agent, regardless of the specific entry times. This is because the later entry of an agent only shifts the time at which it gets ostracized, but will not change the fact that it ever gets ostracized. Since the probability that each agent is ostracized is not affected, the set of stable networks and the probability that each stable network emerges does not change either. Thus the same probability distribution over stable networks as in Corollary 1 will result. □

Proof of Theorem 7

Proof. We prove using an example, shown in Figure A.4. In this network of three agents, suppose that the white agent's expected quality is very high. Suppose both the red (large circle and bolded line) and the blue (small circle and thin line) agents expected qualities are very close to c. Since the white agent's expected quality is very high, the social welfare of the network will be completely determined by the amount of time the blue agent connects with the white agent. By delaying the entry of the red agent, the blue agent is able to stay connected for longer in each realization, and so social welfare increases. □

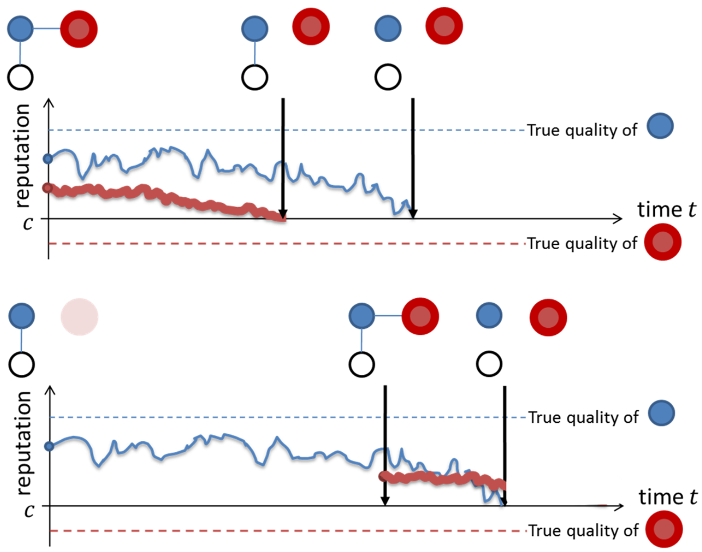

Figure A.4: Example for Theorem 7

[1] J. Scott, Social network analysis, Sage, 2012.

[2] M. O. Jackson, Social and economic networks, Princeton University Press, 2010.

[3] S. Erol, Network hazard and bailouts, Tech. rep., working paper (2015).

[4] M. O. Jackson, A. Wolinsky, A strategic model of social and economic networks, Journal of economic theory 71 (1) (1996) 44–74.

[5] V. Bala, S. Goyal, A noncooperative model of network formation, Econometrica 68 (5) (2000) 1181–1229.

[6] A. Watts, A dynamic model of network formation, Games and Economic Behavior 34 (2) (2001) 331–341.

[7] A. Galeotti, S. Goyal, J. Kamphorst, Network formation with heterogeneous players, Games and Economic Behavior 54 (2) (2006) 353–372.

[8] B. Golub, M. O. Jackson, Naive learning in social networks and the wisdom of crowds, American Economic Journal: Microeconomics (2010) 112–149.

[9] D. Acemoglu, M. A. Dahleh, I. Lobel, A. Ozdaglar, Bayesian learning in social networks, The Review of Economic Studies 78 (4) (2011) 1201–1236.

[10] B. Golub, M. O. Jackson, How homophily affects the speed of learning and best-response dynamics, The Quarterly Journal of Economics 127 (3) (2012) 1287–1338.

[11] F. Vega-Redondo, Building up social capital in a changing world, Journal of Economic Dynamics and Control 30 (11) (2006) 2305–2338.

[12] Y. Song, M. van der Schaar, Dynamic network formation with incomplete information, Economic Theory 59 (2) (2015) 301–331.

[13] M. van der Schaar, S. Zhang, From acquaintances to friends: Homophily and learning in networks, IEEE Journal on Selected Areas in Communications.

[14] D. Li, N. Schürhoff, Dealer networks.

[15] B. Hollifield, A. Neklyudov, C. S. Spatt, Bid-ask spreads, trading networks and the pricing of securitizations: 144a vs. registered securitizations, Trading Networks and the Pricing of Securitizations: 144a vs. Registered Securitizations (August 2014).

[16] B. Chang, S. Zhang, Endogenous market making and network formation, Available at SSRN 2600242.

[17] M. Farboodi, Intermediation and voluntary exposure to counterparty risk.

[18] A. Neklyudov, B. Sambalaibat, Endogenous specialization and dealer networks, Available at SSRN 2676116.

[19] A. Babus, T.-W. Hu, Endogenous intermediation in over-the-counter markets.

[20] C. Wang, Core-periphery trading networks, Available at SSRN 2747117.

[21] F. Blasques, F. Bräuning, I. Van Lelyveld, A dynamic network model of the unsecured interbank lending market.

[22] K. D. Williams, K. L. Sommer, Social ostracism by coworkers: Does rejection lead to loafing or compensation?, Personality and Social Psychology Bulletin 23 (7) (1997) 693–706.

[23] E. D. Wesselmann, J. H. Wirth, J. B. Pryor, G. D. Reeder, K. D. Williams, When do we ostracize?, Social Psychological and Personality Science 4 (1) (2013) 108–115.

[24] S. L. Robinson, J. OReilly, W. Wang, Invisible at work an integrated model of workplace ostracism, Journal of Management (2012) 0149206312466141.

[25] R. Jacobson, D. Aaker, Myopic management behavior with efficient, but imperfect, financial markets: A comparison of information asymmetries in the us and japan, Journal of Accounting and Economics 16 (4) (1993) 383–405.

[26] N. Mizik, The theory and practice of myopic management, Journal of Marketing Research 47 (4) (2010) 594–611.

[27] M. van der Schaar, S. Zhang, A dynamic model of certification and reputation, Economic Theory (2014) 1–33.

[28] D. Duffie, H. Zhu, Does a central clearing counterparty reduce counterparty risk?, Review of Asset Pricing Studies 1 (1) (2011) 74–95.

[29] M. O. Jackson, Y. Zenou, Games on networks, Handbook of Game Theory (2014) 95.

[30] L. Wang, K. Pötzelberger, Boundary crossing probability for brownian motion and general boundaries, Journal of Applied Probability (1997) 54–65.