Reputational Learning and Network Dynamics

Simpson Zhang and Mihaela van der Schaar

Abstract

In many real world networks agents are initially unsure of each other's qualities and learn about each other over time via repeated interactions. This paper is the first to provide a methodology for studying the formation of such networks, taking into account that agents differ from each other, that they begin with incomplete information, and that they must learn through observations which connections/links to form and which to break. The network dynamics in our model vary drastically from the dynamics emerging in models of complete information. With incomplete information and learning, agents who provide high benefits will develop high reputations and remain in the network, while agents who provide low benefits will drop in reputation and become ostracized. We show, among many other things, that the information to which agents have access and the speed at which they learn and act can have tremendous impact on the resulting network dynamics. Using our model, we can also compute the *ex ante* social welfare given an arbitrary initial network, which allows us to characterize the socially optimal network structures for different sets of agents. Importantly, we show through examples that the optimal network structure depends sharply on both the initial beliefs of the agents, as well as the rate of learning by the agents.

I. INTRODUCTION

Much of society is organized in networks and networks are important because individuals typically interact largely or perhaps entirely with those to whom they are connected. A large network science literature studies networks that have already formed (see e.g. Watts and Strogatz (1998)). A smaller microeconomics literature studies[1] the formation of networks - but makes strong assumptions (e.g., homogeneous agents/entities, complete information about

Document Date: April 2014

The authors are indebted to Jie Xu, William Zame, and Moritz Meyer-ter-Vehn for valuable assistance. This paper also benefited from discussions with Matt Jackson, Asuman Ozdaglar, Randy Berry, Robert Kleinberg, and seminar participants at UCLA and Caltech.

Simpson Zhang: simpsonzhang@ucla.edu, Mihaela van der Schaar: mihaela@ee.ucla.edu

[1]See the overview in Jackson (2010) for instance.

other agents). Neither the network science literature nor the microeconomics literature take into account that agents behave strategically in deciding what links to form/maintain/break and begin with incomplete information about others (i.e., they must learn about others). As a result, neither network science nor microeconomics provides a complete framework for understanding, predicting and guiding the formation (and evolution) of real networks and the consequences of network formation.

The overarching goal of this research paper is to develop such a framework. An essential part of the research agenda is driven by the understanding that individuals in a network are not all the same - some workers are more productive than others, some friends are more helpful than others, and some contacts are more useful than others. Furthermore, these characteristics are not known in advance but must be learned over time via repeated interactions. The impact of agent heterogeneity on network structure and dynamics has already been analyzed in several papers from the networks literature, such as Galeotti et al (2006) which shows that heterogeneity of benefits has a strong impact on the connectedness of a network, and Goeree et al (2009) which finds experimentally that large differences in the network formation process will result from heterogeneous agents. However, these previous studies consider only *complete information* when analyzing the network formation process. But complete information seems unrealistic in many real life situations, especially when agents are meeting for the first time. In such circumstances, the presence of incomplete information is crucial and has a huge impact on the network dynamics.

For instance, consider a group of individuals interacting with each other in an online expertise network, such as a forum or interest group. Agents provide benefits to each other by posting advice on a topic, which is publicly observable. These agents are likely to be initially uncertain of the benefits provided by other agents with whom they have little experience, and linking entails costs in terms of time, energy etc., so individuals will only interact with another individual (over time) if the value derived from that agent is high. As time progresses, individuals *observe* the value of the advice produced by other individuals, *update* their belief about the value of each agent, and *change* their linking decisions as a result. Thus agent learning causes the network topology to evolve over time. But the network topology also impacts agent learning, as agents with more links are giving out more advice and thus revealing information about themselves more quickly. Therefore the learning and the network topology are closely intertwined and co-evolve.

Our model takes into account that agents behave strategically and that they begin with incomplete information about each other and thus, must learn which connections to form and maintain and which to break through their continued interactions. We consider a continuous time model with a group of agents who are linked according to a network and who send flow benefits to their neighbors. These flow benefits are noisy and evolve according to a Brownian motion with drift equal to each agent's true quality level, and variance determined by the amount of current neighbors[2]. Agents observe all the benefits that their neighbors produce, and update their beliefs about their neighbor's quality via Bayes rule. They will maintain links with neighbors that provide high benefits, but will cut off links with neighbors that provide low benefits. Thus the network will evolve over time as agents learn about each other and update their beliefs. Since the amount of links an agent has determines the rate of learning about that agent, the rate of learning changes as the network changes, leading to a co-evolution of the network topology and agent learning. Our model is highly tractable and allows us to completely characterize network dynamics and give explicit probabilities of the network evolving into various configurations. In addition, we are able to describe the entire set of stable networks, and analytically compute the probability that any single stable network emerges.

We also study the implications that learning has on social welfare and the efficiency of a network. Our results show that learning has a beneficial aspect: agents that are of low quality are likely to produce low signals and will eventually be forced out of the network. Learning also has a harmful aspect: even high quality agents 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 neighbor slightly, but it also receives a large benefit if its neighbor is of very high quality. Thus 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, and not the benefit they are providing their neighbors. Thus, there is a negative externality every time a link is severed. In many situations, the 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 agents 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 agents is good, because more information will be revealed about that marginal quality

[2]The more neighbors an agent has, the more information the agent reveals about itself and so the smaller the variance.

agent, allowing its neighbors to more quickly sever their links to it. Thus the impact of learning can be either positive or negative depending on the specific network topology.

Our welfare results also have important implications for network planning and may be useful, for example, in guiding the formation of networks by policies of a human resources department, online community, etc. Due to the varying effects of learning, we show that the optimal network structure will be quite different for different groups of agents. For instance, when agents all have high qualities, the optimal network will be fully connected (which allows all agents to benefit fully from their repeated interactions). On the other hand, if some agents have low initial reputations, then a fully connected network may not be optimal because it will be desirable to isolate low quality agents (or clusters of low quality agents). Because of the negative effects of learning, it may be optimal to prevent two agents from linking with each other, even if such agents have initial expected qualities higher than the linking cost. If such agents did decide to link, they would both be sending more information about themselves through this link. Thus each agent, as well as the overall network, may become worse off through the formation of this link due to the faster learning caused by the link. In some cases, a star or a core-periphery network 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 very high expected quality and the other group with moderate expected quality. By placing the high quality agents in the core and the moderate quality agents in the periphery, the high quality agents are able to produce large benefits for the network, and the potential harm from the moderate quality agents is minimized.

II. Literature Review

Our paper represents a novel contribution to the network formation literature, by being the first to consider asymmetric 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 network topologies. Other papers in the network literature have usually studied network dynamics only in settings of complete information when agents perfectly know each other's qualities. For example, the papers by Jackson and Wolinsky (1996), Bala and Goyal (2000), Watts (2001), and Galeotti and Goyal (2010) all consider networks where the agents have complete information. In these models, agents are aware of all 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. As one link is severed or formed, the benefits produced by other links changes as well, which causes other links to sever or form as well in a chain reaction. For some networks, such as informational networks these indirect benefits seem important, as an agent who has many neighbors will likely produce higher quality information as well. However, in other networks such as collaboration networks these indirect benefits are less relevant and it is the quality of each specific agent that matters the most. We argue that the network dynamics that arise in such situations are much more due to incomplete information and learning than due to changes in the value of indirect benefits.

Contrary to the other papers, we do not assume any indirect benefits in our model and focus instead on the dynamics resulting from the incomplete information and learning by the agents. Agent learning has a strong influence on the network formation process in a way that could not arise without incomplete 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 thus has a strong effect 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, Watts (2001) 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 any 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. Thus 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.

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 in our model and the choice of Brownian motion as the information process, both of which allow for closed form equations of the probabilities that different networks emerge. In contrast other networks papers such as Jackson and Wolinsky (1996) and Bala and Goyal (2000) use discrete time models that do not allow for such clean closed form expressions. While the 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 lets us compare the *ex ante* optimality of different initial network structures, as well as provide general results for when certain networks 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 will be optimal if the agent's initial expected qualities are all higher than the cost of maintaining the links. But when learning is at an intermediate rate, it may be optimal not to have all agents connected with each other even if their expected qualities are higher than the linking cost, due to the externalities associated with learning. 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.

III. Model

We consider an infinite horizon continuous time model with a finite number of agents denoted by $V = \{1, 2, ..., N\}$. The agents are able to form links with each other, and these links are undirected and require mutual consent. There is an underlying network constraint $\Omega = \{\omega_{ij}\}$ that specifies which pairs of agents are able to link with each other. 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. by the human resources department in a company for a network of employees. 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 $k_i^t = \sum\limits_j g_{ij}^t$ be the number of links that agent i has 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 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 Brownian motion $db_{ij}(t) = q_i dt + \tau_i^{-1/2} dZ_j(t)$, where the drift is the true quality q_i and the variance depends on τ_i, which is denoted as the base precision of agent i[3]. $Z_j(t)$ is a standard zero-drift and unit variance Brownian motion, and represents the random fluctuations in the benefits of each interaction. $Z_j(t)$ is independent over all j, so the individual benefits 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[4].

Since we have assumed a Brownian motion 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 Brownian motion $dB_i(t) = q_i dt + (k_i^t \tau_i)^{-1/2} dZ(t)$ 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 $Z(t)$ 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 that agent's quality. Therefore the topology of the network strongly affects the rate of learning about agent quality.

For each agent i, we define the agent's benefit history as the history of all previous benefits, $\mathcal{H}_i^t = \{b_{ij}^{t'}\}_{t'=0}^{t}$. If at time t all 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 hence, the Brownian motion of agent i is stopped at the current level. As mentioned, there is a prior belief of an agent i's quality $\mathcal{N}(\mu_i, \sigma_i^2)$ and agents will update 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

[3]We can think of the base 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 or chance encounters), or factors like the agent's personality.

[4]This is an important assumption to maintain the tractability of the model. It can be interpreted, for instance, through 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 other neighbors of i, so that the neighbors maintain the same beliefs.

also normally distributed[5]. 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 assume that agents are myopic, and thus consider only the current flow benefit when making linking decisions[6]. Therefore, in the basic model, agent i will cut off the 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.

IV. Model Dynamics and Stability

A. *Model Dynamics*

The dynamics of the model will 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 assume 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 does not need to be considered). Thus we will have $G^0 = \Omega$. 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 other agents 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

[5]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.

[6]Such an assumption is common within the networks literature to maintain tractability, see Jackson and Wolinsky (1996) or Watts (2001) for instance.

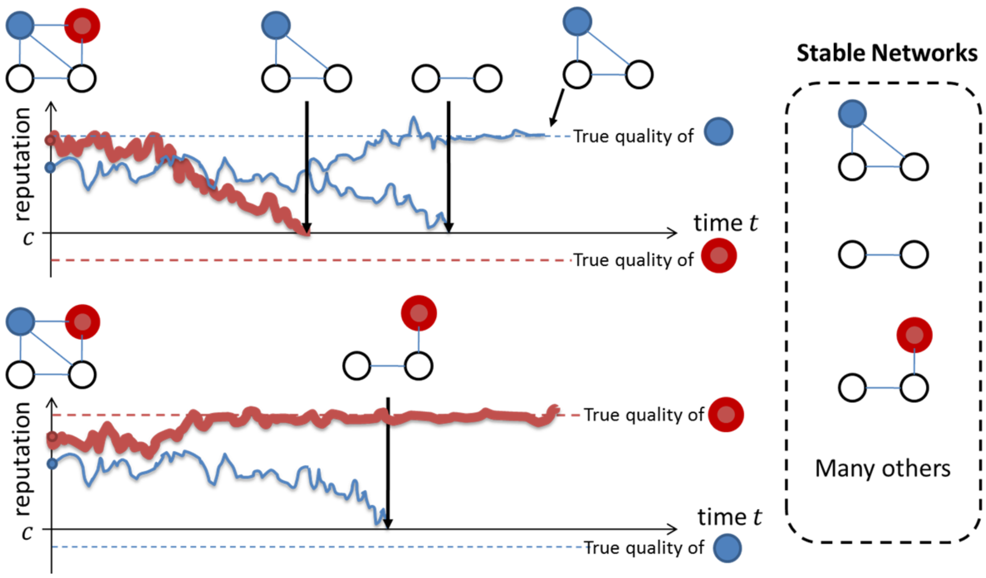

Fig. 1. Illustration of Possible Network Dynamics: From the same initial reputations for the red and blue agents, many different network dynamics and stable networks are possible. In the top graph the red agent has a true quality less than c and so will be ostracized from the network for certain at some time, while the blue agent has a true quality above c and so may or may not be ostracized from the network depending on the signals it sends. Each event leads to a different stable network, one with and one without the blue agent. In the bottom graph it is the blue agent who has a true quality lower than c and so will be ostracized for sure, whereas the red agent could potentially stay in the network indefinitely.

note that many different stable networks could potentially emerge depending on the true qualities of the agents and the signals they produce. Figure 1 shows the different network dynamics that could emerge even if the initial reputations of the agents are fixed, due to the uncertainty about the true qualities of the agents as well as the randomness in the signals they send.

B. *Stable Networks*

As mentioned, we call the limiting network structure when 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. We can characterize the set of stable networks that can emerge given a set of agents and a network constraint Ω, as well as how the rate of learning affects the set of 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. 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 done 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.

Proposition 1. *$P(S_i)$ depends only on the initial quality distribution and the link cost and can be computed by*

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

Proof. See appendix. □

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 or excluded in the stable network. Also note that the probability an agent is part of the stable network is independent of that agent's base precision τ_i. Thus 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. 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 its initial mean quality μ_i, decreasing in the variance of its initial quality σ_i^2, and decreasing in the link cost c. Moreover,* $\lim_{\mu_i \to \infty} P(S_i) = 1$, $\lim_{\sigma_i \to 0} P(S_i) = 1$, $\lim_{c \to -\infty} P(S_i) = 1$.

Proof. See appendix. □

These properties are intuitive since an agent with a higher mean quality and smaller variance is less likely to have its reputation drop below c, and so is less likely to become ostracized. Moreover, lowering the linking cost also reduces the hitting probability since the agent's reputation would now have to fall lower to be excluded from the network.

As mentioned, G^∞ must be a subset of G^0. Further, it can contain links only amongst pairs of agents that are both included in the stable network and were 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[7].

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 on the aggregate level either. This immediately implies the following theorem. We can also use the characterization of stable networks along with Proposition 1 to derive an analytic expression of the probability that any specific stable network emerges, which is presented in the corollary below. Figure 4 in the appendix shows an application of the corollary to a network of three agents.

[7]For instance suppose that the network comprises only two agents i and j. Then the event in which only S_i occurs and the event in which only S_j occurs lead to the same stable network structure.

Theorem 1. *The base 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.*

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 $g_{ij} = I\{S_i\}I\{S_j\}I\{g_{ij}^0 = 1\}$. 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. □

Corollary 2. *The probability that a network G is a stable network is given by $\sum\limits_{\{S_i\}} \prod_i P(S_i)$ 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.

V. Welfare Computation

We will analyze social welfare from an *ex ante* perspective, given only the network constraint Ω and the prior distributions of the agent qualities. Importantly the *ex ante* welfare is calculated before the agent qualities are known and any signals are sent. This type of welfare is the most suitable for the design settings we will consider later, as it requires the least knowledge on the part of the network designer. Denote $P(L_{ij}^t|q, G^0)$ as 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. We can define the *ex ante* social welfare W formally as as follows

$$W = \int_{q_1=-\infty}^{\infty} \cdots \int_{q_N=-\infty}^{\infty} \sum_{i,j} \int_0^{\infty} e^{-\rho t}(q_j - c)P(L_{ij}^t|q, G^0)dt\phi(\frac{q_N-\mu_N}{\sigma_N})dq_N \ldots \phi(\frac{q_1-\mu_1}{\sigma_1})dq_1 \tag{2}$$

Directly computing the $P(L_{ij}^t|q, G^0)$ in this formula is very difficult since the linking decisions of agents i and j are coupled with the linking decisions of other agents in the network. For instance, if a neighbor of agent i is ostracized, i would then have fewer neighbors and thus sends information more slowly, affecting the speed at which j learns about it, and thus the speed at which j updates its linking decision. We will not try to compute this probability directly, but instead take an indirect approach. This approach utilizes the fact that the *ex ante* social welfare is an expectation over all the possible *ex post* signal realizations. Thus we can calculate the *ex ante* welfare by integrating over all possible realizations of the *ex post* welfare, which is a much more tractable term.

A. *Ex post welfare*

Consider an *ex post* realization of hitting times $\varepsilon = \{\varepsilon_i^{t_i}\}_{i\in\mathcal{V}}$ 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

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

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

and we have

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

where $P(S_i)$ can be explicitly computed according to Lemma 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 thus 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 agent. We can calculate the *ex post* surplus that an agent i receives given ε as

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

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

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

Summing over all agents, the social welfare given the *ex post* realization ε is therefore

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

Taking the expectation over the events ε, the *ex ante* social welfare is thus $W = E_\varepsilon[W(\varepsilon)]$. In order to compute the *ex ante* social welfare, we still need to know the distribution of the t_i which again 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.

B. 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. Thus 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 under its base precision, and then scale this time downwards proportionately based on the network effect. 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 $\hat{\varepsilon}_i^{t_i}$ can be explicitly computed in the following lemma.

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

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

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

Proof. See appendix. □

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 time may vary for each $\hat{\varepsilon}$. Let $M : [0,\infty]^N \to [0,\infty]^N$ 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 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. 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 summarizes our result.

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

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

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

VI. Impact of Information and Learning

In this section, we study the impact of learning on the *ex ante* welfare given an initial network G^0. In particular, we will show how the Brownian motion base precisions change each individual agent's 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. Thus the *ex ante* social welfare without learning can be computed by

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

A. Overall Impact of learning

Let $W(\tau_1, ..., \tau_N)$ be the social welfare when agents learn each others' true quality with the base precisions being $\tau_1, ..., \tau_N$. The next theorem states that learning has a negative impact on the social welfare overall.

Theorem 3. $W(\tau_1, ..., \tau_N) < W^*$.

Proof. See appendix. □

There are two main factors that are work in this result. First, the myopia of the agents causes the learning to be done inefficiently. Secondly, 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, they lead to a reduction in social welfare overall. More precisely, when a link l_{ij} is severed due to agent j's reputation hitting c, agent i does not lose 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[8]. 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.

[8]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 Schaar and Zhang (2014). In Section VIII we discuss a possible solution for this problem.

B. Impact of information precision

In this subsection, we investigate how the speed of information sending (i.e. the base precision) affects the individual welfare. The faster an agent sends information about its own reputation, the faster the other agents will learn its true quality (if the link is not broken).

The next proposition characterizes the impact on the agent's own welfare.

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

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 hitting time realization of all agents. Thus agent i's welfare $W_i(\varepsilon)$ is not affected. If $t_i < \infty$, then the welfare of agent i depends on (1) the mean 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'. □

This result shows that sending information about itself is bad for the agent's own welfare. This is because in each realization that the agent is ostracized from the network, the agent will now be ostracized sooner and hence enjoys less benefits from others.

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

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

Proof. Consider any *ex post* realization $\varepsilon = \{\varepsilon_i^{t_i}\}_{i \in V}$. If $t_i = \infty$, then increasing agent i's base 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 Theorem 1, the new hitting time t_i' is sooner if agent i's base 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 base 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. □

Since the information sending speed between times t_i' and t_i becomes slower for agent j, agent j's hitting time is larger. Thus, agent j would enjoy more benefit from its neighbors. We can do a similar analysis for more distant agents when the two agents are connected through a unique path. This is summarized in the corollary 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 base 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 thus, the impact of any agent's information sending speed on any other agent's welfare can be completely characterized. However, when there are multiple paths between agents, which implies there are cycles in the network, the impact of learning speed of an agent on the other agents' welfare is much less clear. The following proposition shows that even for the immediate neighbor, the impact could be totally opposite 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 base precision decreases its immediate neighbor's welfare.*

Proof. We prove by constructing a counterexample, which is shown in Figure 2. 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 quality is perfectly known 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 but the variance is small (e.g. on the order of $O(1/K^2)$. Hence, agent 2 almost

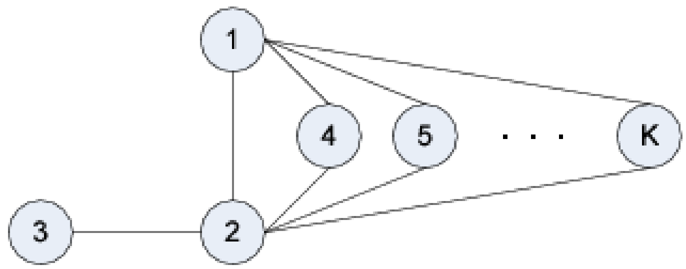

Fig. 2. Counterexample for Proposition 4

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 base 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 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. $\square$

We have seen that increasing the information sending speed of agent i can be both good or bad for other agents depending on their locations in the network and their relation with agent i. Now we study how it affects the social welfare as a whole.

Theorem 4. *Increasing all agents' base precision by the same ratio $\lambda > 1$ decreases the social welfare. The social welfare is decreased to no less than $1/\lambda$ of the original social welfare.*

Proof. Consider any *ex post* realization ε. Since the base precision of all agents are increased by the same ratio λ, the hitting time of all agents will be decreased by the same ratio $1/\lambda$. Due to discounting, the welfare of each agent is reduced to no less than $1/\lambda$ of the original welfare. Thus, for each agent, it obtains less benefit from its neighbors due to the earlier hitting time. Because this is true for all *ex post* realization, the *ex ante* social welfare also decreases to no less than $1/\lambda$ of the original social welfare. $\square$

Theorem 4 proves a negative effect of the overall learning speed on the social welfare[9]. This negative effect of learning is partly due to agent myopia, which results in the system failing to internalize the benefits due to learning. A potential method of addressing the negative effect of learning is to provide a subsidy to the agents for each link, which will cause the agents to experiment for longer. We will discuss this extension in Section VIII.

Because of the negative effect of learning, it may be optimal to prevent two agents from linking with each other, even if such agents have initial expected qualities higher than the linking cost. If such agents did decide to link, then they would increase the rate at which the other sends information. Thus, it is important to carefully consider which agents are connected with which other agents when designing a network. This is the topic of the next section.

VII. Optimal Initial Networks

In this section, we study which initial networks G^0 maximize the social welfare. Equivalently, we could think of designing the network constraint Ω to maximize social welfare, by choosing which agents are able to link with which other agents.

A. *Fully connected networks*

It is natural to think that a fully connected network could be the optimal initial network 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. Hence, if the designer is completely impatient (i.e. $\rho \to \infty$), then having all agents fully connected generates the highest social welfare, since the designer cares only about the current time period, and when time is short almost no new information can be learned. Surprisingly though, the fully connected network is also optimal on the other extreme, when the designer is completely patient (i.e. $\rho \to 0$). Note also that the level of patience is closely connected to 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.

[9]Note however that increasing the base precision of only a single agent instead of all agents simultaneously can actually increase welfare. Proposition 4 shows that the agent's neighbor's welfare could increase, and this increase can outweigh the loss in welfare of the agent itself.

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 initial network is the fully connected network.*
2. Fix the other parameters of the model and suppose the agents' base 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 initial network is the fully connected network.

Proof. See appendix. □

When the designer is either completely impatient or patient, 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. 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 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 optimal in the case where the agents are homogeneous and have very high initial qualities.

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

Proof. Since $\lim\limits_{\mu_i \to \infty} P(S_i) = 1$, when $\mu_i \to \infty$, the optimal initial network is the fully connected network because all agents will be in the stable network with probability 1. Hence there must exist a sufficiently large $\bar{\mu}$ such that the social welfare is dominated by the stable network that includes all agents. □

B. *Core-periphery networks*

As agents become more heterogeneous in terms of their initial mean quality, different network structures other than the fully connected network can emerge to be the optimal initial network. Suppose agents are divided into two 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 initial network has a core-periphery structure.

Theorem 5. *Suppose that there are two groups of agents, one with initial quality μ_L and one with initial quality μ_H, and suppose that the variances of both groups are a constant σ_L^2 and σ_H^2. Fixing all other parameters, there exists $\bar{\mu}$ such that $\forall \mu_H > \bar{\mu}$, the optimal initial network is a core-periphery network where all high type agents are connected with all other agents and no two low type agents are connected.*

Proof. We first show that all high type agents should connect to all other agents. This is based on a similar argument as in Proposition 6. Since when $\mu_H \to \infty$, all high type agents will stay in the stable network, the welfare a high type agent obtains from other agents is maximized when the high type agent is connected to all other agents in the initial network. 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 agents in the initial network.

Next we show that all low type agents do not connect to each other. When $\mu_H \to \infty$, the welfare obtained by any low type agent j is dominated by that from its links with high type agents. Having additional links with other low-type agents reduces the hitting time of agent j in the event that it gets ostracized, thereby reducing agent j's welfare. Therefore, low type agents do not connect to each other in the optimal initial network. □

C. *Ring networks*

In this section we focus on a special type of network: a ring network. Suppose for convenience that agents are homogeneous in terms of initial mean quality and variance. Assume that under the network constraint Ω each agent has at most have two neighbors. Hence, for a given number of agents, they would only be able to form one or multiple rings of different sizes. We study how the size of different rings affects the welfare an agent obtains and hence, we can determine the optimal size of rings that agents should form together. Let $W(n)$ denote the welfare an agent can obtain if it is in a ring of size n under the network constraint Ω.

Proposition 7. *The optimal size of a ring network is 3 agents.*

Proof. Consider a ring network consisting of three agents i, j, k and we focus on the welfare of agent i. Agent i obtains a positive benefit in two cases: (1) realizations in which both agents j and k's reputation never hit c; (2) realizations in which exactly one of agents j and k's reputation never hits c. The probabilities that these two cases happen are independent of the network structure by Proposition 1. In the first case, having additional agent(s) between agent j and k does not affect agent i's realization and hence, agent i's welfare is not affected. In the second case, having additional agent(s) between agent j and k will change i's realization with positive probability. Consider a realization in which agent k's reputation never hits c and agent j's reputation hits c at t_j. In the ring of size 3, agent j's direct neighbor besides i (i.e. agent k) never hits c. When there are additional agents, it is either the case that agent j's new direct neighbor never hits c or hits c before infinity. If agent j's new direct neighbor hits c before infinity, then agent j's new hitting time may increase and hence agent i's new hitting time may decrease, leading to a lower welfare for agent i. □

Corollary 4. *If n is odd, then $W(n) > W(m), \forall m > n$. If n is even, then $W(n) < W(m), \forall m > n$.*

Proof. The proof is similar to that of Proposition 7 except we take into account the odd-even effect discussed in Corollary 3. □

D. Star Networks

In star networks, there are also important facts to consider. The network is minimal like the ring network, and depends very much on the central agent, because that agent is connected with all other agents and is generating very high benefits for these other agents. The central agent is thus the most important agent to consider, and choosing the best agent to be in the center is crucial to the welfare of the network. There are two facts to consider, the initial mean and the variance of the agent. A high initial mean is clearly beneficial, because it increases the flow benefits that all the other agents will receive who are connected to the central agent. However, a higher initial variance is harmful because it causes the link to be more easily dissolved, and thus social welfare would be lowered greatly. Therefore there is a trade off between the initial mean and the variance of the central agent: we want a higher mean but a lower variance. In particular, choosing the agent with the highest mean may not be optimal because if that agent also has a higher variance.

In addition the information speed of the central agent matters greatly. We would like the central agent to send information more slowly, because that will increase the stability of the links in the network. Thus we want to choose a central agent with a low information precision as well. These are the various tradeoffs that need to be considered when thinking about the best central agent.

VIII. 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 three possible extensions that could alleviate this issue and allow for higher social welfare.

A. 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 may wish to give workers awards or bonuses for collaborating with colleagues. Formally, for every link that an agent maintains, it receives an extra flow benefit of δ from the network designer. This linking subsidy would not matter in the social welfare computation since it is just 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 thus causes the agents agents to learn more information about their neighbor's quality and break only if it is really 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 6. *There exists $\bar{\delta}$ such that $\forall \delta > \bar{\delta}$, $W(\delta) > W^*$. Moreover, $\lim_{\delta \to \infty} W(\delta) = W^*$.*

Proof. See appendix. □

Note that by Theorem 3, this result also shows that the social welfare is higher than the standard network model with no subsidy. Thus by imparting subsidies on agents to encourage them to experiment for longer, the social welfare is higher than previously. The intuition is that when the link subsidy is high enough, any link that is broken will involve an agent that was of really bad quality, and so social welfare will improve when this link is broken. Thus 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. Thus having a linking subsidy is beneficial for the network, but the subsidy cannot be set too high either in order to maximize social welfare.

B. *Link formation*

Another way to make learning more socially beneficial is to allow agents to form new links with others whose reputations are high enough. In this extension, we assume that forming a new link with any agent that they are not connected to under the network constraint Ω incurs 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 represent for instance the cost of time and energy in associating with a new agent, or some explicit cost that the network designer is imposing for two agents that want to link.

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 quality agents to link together will improve social welfare due to the high mutual benefits they 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 7. *There exists $\bar{\gamma}$ such that $\forall \gamma \geq \bar{\gamma}$, $W(\gamma) > W$.*

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, which 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' reputation 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 change is bounded, i.e. $W'(\varepsilon|\zeta_2) - W(\varepsilon|\zeta_2) \geq -e^{-\rho t^*}B$ since there are a finite number of agents. When ζ_1 occurs, forming a link between agent i and j does not change the welfare of other agents but only increases the welfare of agent i and j by at least $\frac{e^{-\rho t^*}}{\rho}\frac{(c+\bar{\gamma})-(1-P(\zeta_1))c}{P(\zeta_1)}$. Hence the change in social welfare satisfies

$$W'(\varepsilon) - W(\varepsilon) \geq P(\zeta_1)\frac{e^{-\rho t^*}}{\rho}\frac{(c+\bar{\gamma})-(1-P(\zeta_1))c}{P(\zeta_1)} - e^{-\rho t^*}B \tag{14}$$

$$\geq e^{-\rho t^*}[\frac{(c+\bar{\gamma})-c}{\rho} - B] \tag{15}$$

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

□

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 high reputations. Thus the social welfare generated by a new link is likely to be 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 if there are a group of moderate quality agents that are all linked to a very high quality agent, but separated from each other according to Ω. This is similar to the situation considered in Theorem 5. In such a case, allowing the moderate quality agents to link at too low of a reputation 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. Thus allowing for extra link formation can improve welfare, but the threshold for the link being formed must be sufficiently high as well.

C. Agent Entry

Our model can also be tractably extended to allow agents to over time into the network, instead of all being present in the network at the beginning. 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 who have $e_i > 0$ enter later on. These entry times are fixed and known to the agents in the model. There is still the network constraint Ω over the set of all N potential agents that 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, any new agent that arrives will want to be connected to by its neighbors. However, the new agent would not want to connect with one of its neighbors if that neighbor has already been ostracized previously within the network. Thus 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 change the benefits and amount of information produced by the network.

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

Proposition 8. *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.*

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. □

Although the set and distribution of stable networks is not affected by agent entry, the social welfare will be affected. We note that we can 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 on 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 benefit social welfare due to the increased benefits that the agent and its neighbor are able to extract from the network. In total, the second effect may dominate the first, so that delaying entry for an agent raises social welfare overall.

Theorem 8. *Fixing the other parameters of the network, increasing an agent's entry time* e_i *can increase social welfare.*

Proof. We prove using an example, shown in Figure 3. In this network of three agents, suppose that the gold agent's quality is very high and known, so there is no uncertainty about this agent. Suppose both the red and the blue agents qualities are unknown and close to c. Since the gold agent's 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 gold 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. □

In this example, note that although delaying the entry of the red agent is helpful, it is still better to have the red agent enter at some finite time instead of never entering. This is because the blue agent's reputation will eventually converge to its true quality by the law of large numbers, and in the case where the blue agent has a good quality, enabling a link with the red agent will produce positive benefits. In addition, after waiting for a sufficiently large amount of time, the probability that the blue agent ever becomes ostracized if it hasn't already goes to zero, so the red agent is unlikely to impact the blue agent's connection with the gold agent. Thus delaying the entry of the red agent is beneficial, but the red agent should not be excluded from the network altogether. This argument holds in general for any network as well, all agents should have finite entry times to ensure optimal social welfare.

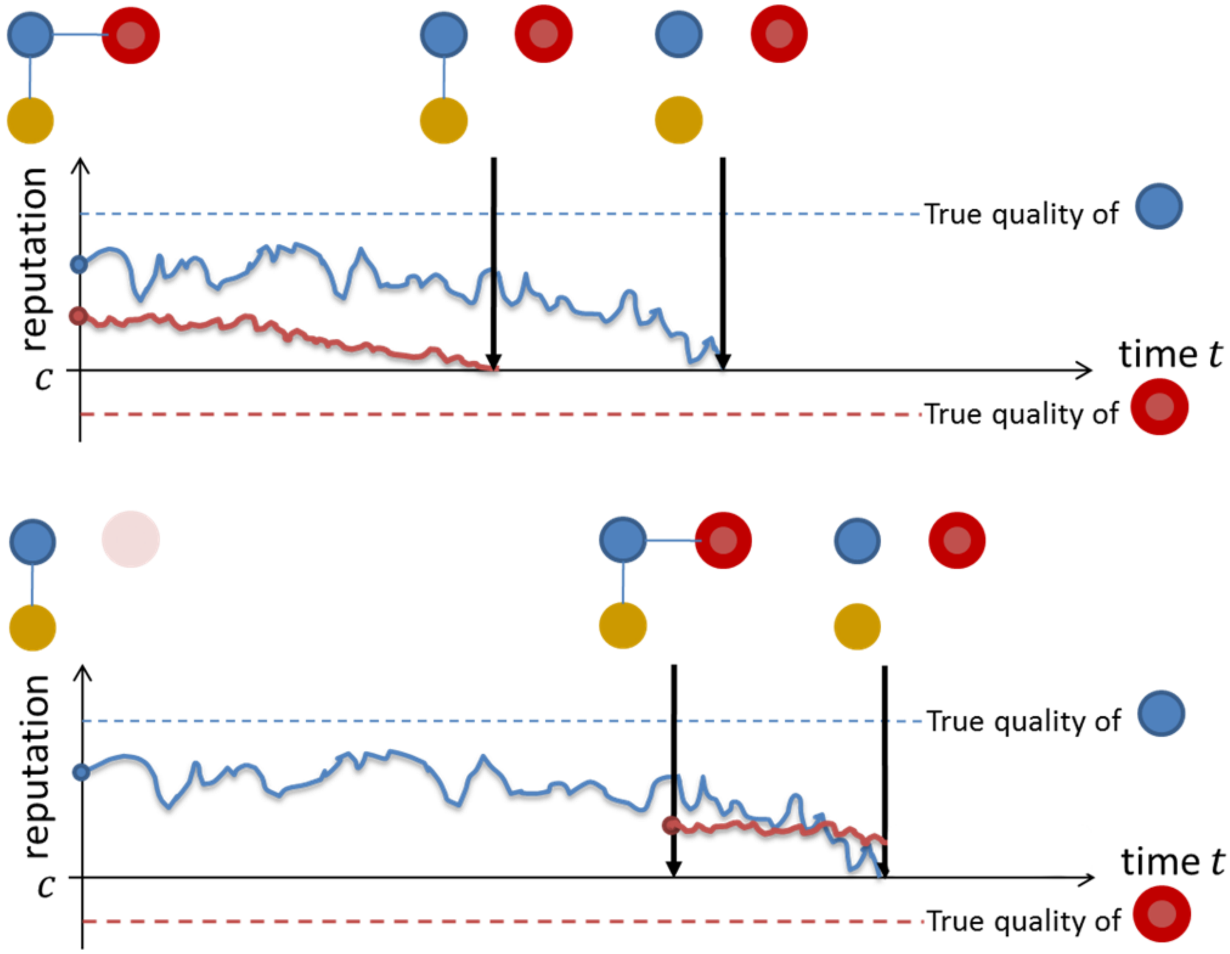

Fig. 3. Example for Theorem 8

IX. 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 the social welfare depending on the specific network topology, and analyzed what optimal 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.

Our results could be extended in several interesting ways. One natural extension would be to allow the qualities of agents to change over time. In the simplest extension, the agent's quality

q_i itself would evolve according to a stochastic process, for instance a Brownian motion. More interestingly, it would be natural to assume that the evolution of quality depends endogenously on the information the agent receives so that agents who receive better information tend to have higher quality and hence to generate better information in the future. Thus, the structure of the network and the quality of the agents in the network co-evolve. Higher quality agents may link to agents that are also of higher quality, and so their qualities would improve quickly, while lower quality agents may struggle to find good agents to link with, and their qualities would decline as a result.

APPENDIX

PROOF OF PROPOSITION 1

Proof. Suppose for now that agent i's reputation always evolves at the constant base 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 Wang and Pötzelberger (1997)) and is given by

$$P(S_i^t|q_i) = \Phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{16}$$

$$-\exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c))\Phi\left(\sqrt{t\tau_i}(q_i - c) - \frac{\frac{1}{\sigma_i}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{17}$$

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

$$P(S_i|q_i) = \lim_{t\to\infty} P(S_i^t|q_i) \tag{18}$$

- 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 $q_i < c$, $\Phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i-c)}{\sqrt{t\tau_i}}\right) \to 0$ 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

$$P(S_i) = \int_c^\infty (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i \tag{19}$$

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 base precision τ_i. □

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

$$P(S_i) = \int_{c-\mu_i}^\infty (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(\mu_i - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx \tag{20}$$

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

$$P(S_i|\mu_i') = \int_{c-\mu_i'}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i' - c)(\mu_i' - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx \tag{21}$$

$$> \int_{c-\mu_i}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i' - c)(\mu_i' - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx \tag{22}$$

$$> \int_{c-\mu_i}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(\mu_i - c + x)))\phi\left(\frac{x}{\sigma_i}\right) dx = P(S_i|\mu_i) \tag{23}$$

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

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

$$P(S_i) = \int_{c-\mu_i}^{\infty} \phi(\frac{x}{\sigma_i})dx - \int_{c-\mu_i}^{\infty} e^{-\frac{2}{\sigma_i^2}(\mu_i - c)(\mu_i - c + x))} \frac{1}{\sigma_i\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x}{\sigma_i}\right)^2} dx \tag{24}$$

$$= \int_{c-\mu_i}^{\infty} \phi(\frac{x}{\sigma_i})dx - \int_{c-\mu_i}^{\infty} \frac{1}{\sigma_i\sqrt{2\pi}} e^{-\frac{2}{\sigma_i^2}(2(\mu_i - c) + x)^2)} dx \tag{25}$$

$$= \int_{c-\mu_i}^{\infty} \phi(\frac{x}{\sigma_i})dx - \int_{\mu_i - c}^{\infty} \phi(\frac{x}{\sigma_i})dx = \int_{c-\mu_i}^{\mu_i - c} \phi(\frac{x}{\sigma_i})dx \tag{26}$$

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

$$P(S_i|c) = \int_{c}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i \tag{27}$$

$$< \int_{c}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c')(q_i - c')))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i \tag{28}$$

$$< \int_{c'}^{\infty} (1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c')(q_i - c')))\phi\left((q_i - \mu_i)\frac{1}{\sigma_i}\right) dq_i = P(S_i|c') \tag{29}$$

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 $c' < q_i < c$, $1 - \exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c)) > 0$. □

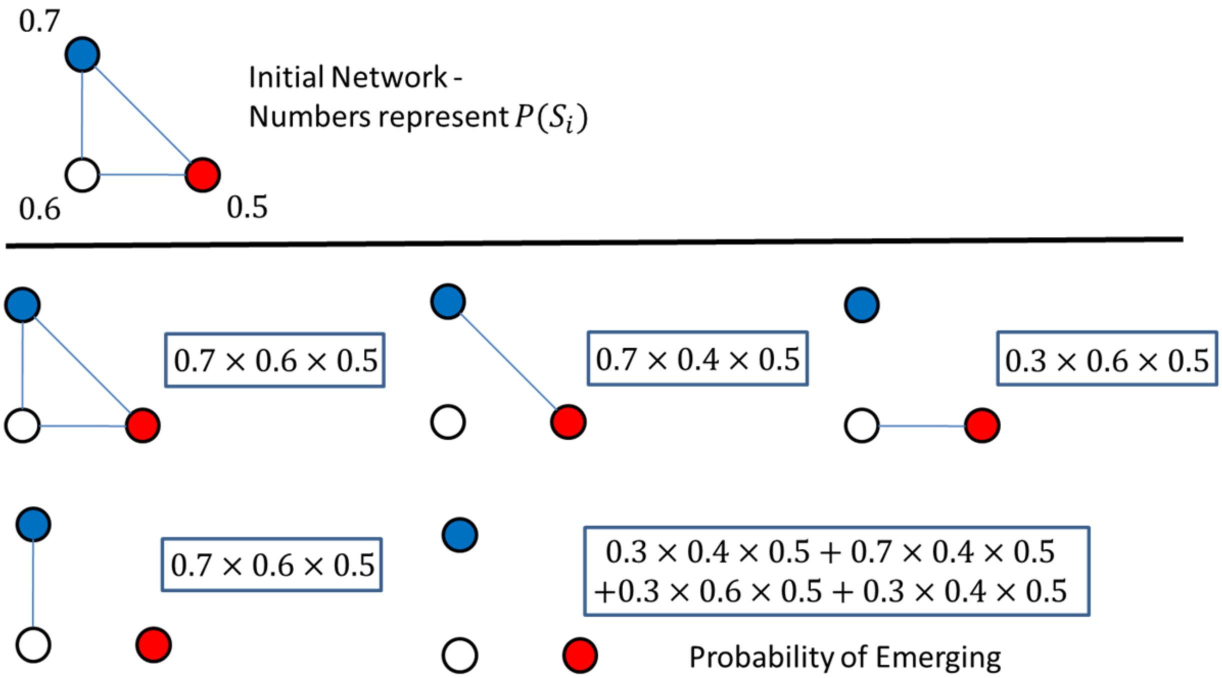

Fig. 4. Set of Stable Networks Given an Initial Network: This figure shows the five possible stable networks that could emerge given an initial network of three agents. In addition, $P(S_i)$ is given for all the agents, which allows us to calculate the exact probability of each of these networks emerging. 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 Lemma 1

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

$$P(S_i^t|q_i) = \Phi\left(\sqrt{t\tau_i}(q_i - c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{30}$$

$$-\exp(-\frac{2}{\sigma_i^2}(\mu_i - c)(q_i - c))\Phi\left(\sqrt{t\tau_i}(q_i - c) - \frac{\frac{1}{\sigma_i}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{31}$$

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

$$f(\hat{\varepsilon}_i^t|q_i) = -\frac{1}{2}\left(\sqrt{\tau_i}(q_i-c)t^{-1/2} - \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i-c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{32}$$

$$+e^{-\frac{2}{\sigma_i^2}(\mu_i-c)(q_i-c)}\frac{1}{2}\left(\sqrt{\tau_i}(q_i-c)t^{-1/2} + \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i-c) - \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{33}$$

$$= -\frac{1}{2}\left(\sqrt{\tau_i}(q_i-c)t^{-1/2} - \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i-c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{34}$$

$$+\frac{1}{2}\left(\sqrt{\tau_i}(q_i-c)t^{-1/2} + \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\right)\phi\left(\sqrt{t\tau_i}(q_i-c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{35}$$

$$= \frac{\mu_i - c}{\sigma_i^2\sqrt{\tau_i}}t^{-3/2}\phi\left(\sqrt{t\tau_i}(q_i-c) + \frac{\frac{1}{\sigma_i^2}(\mu_i - c)}{\sqrt{t\tau_i}}\right) \tag{36}$$

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

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 $G^0_{ii^*} = 1$.

end while

Fig. 5.

ALGORITHM FOR COMPUTING HITTING TIME MAPPING FUNCTION M

PROOF OF THEOREM 3

Proof. The social welfare with learning is given by (2). Consider the *ex ante* social surplus W_{ij} that agent i obtains from the link with agent j. The *ex ante* social welfare is simply the summation over all such i and j. W_{ij} can be computed as

$$W_{ij} = \int_q \int_0^\infty e^{-\rho t} P(L_{ij}^t | q)(q_j - c) dt \phi(q) dq \tag{37}$$

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

$$W_{ij} = \int_q \int_0^\infty e^{-\rho t}(q_j - c) dt \phi(q) dq - E_{t^*}[\int_{t^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t^*) dt] \tag{38}$$

$$= \int_0^\infty e^{-\rho t}(\mu_j - c) dt - E_{t^*}[\int_{t^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t^*) dt] \tag{39}$$

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 $t_j^* = t_j$, namely agent j's reputation hits c before agent c. Then

$$W_{ij} = \int_0^\infty e^{-\rho t}(\mu_j - c) dt - E_{t_i^*}[\int_{t_i^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t_i^*) dt] - E_{t_j^*}[\int_{t_j^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t_j^*) dt] \tag{40}$$

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, $E_{q_j}(q_c - c | t \geq t_i^*) > 0$ because at t_i^* the expected quality of q_i is still no less than c and strictly greater than c with positive probability. Therefore,

$$W_{ij} = W_{ij}^* - E_{t_i^*}[\int_{t_i^*}^\infty e^{-\rho t} E_{q_j}(q_j - c | t \geq t_i^*) dt] < W_{ij}^* \tag{41}$$

Summing over all such pairs of i and j, we conclude that the social welfare with learning is strictly less than that when there is no learning. □

PROOF OF PROPOSITION 5

Proof. If the designer is completely impatient, it only cares about the social surplus at time 0. 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.

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. □

PROOF OF THEOREM 6

Proof. (1) Consider the welfare on link l_{ij}. Let $\zeta_i(t^*)$ denote the event in which agent i's reputation hits $c-\delta$ at time t^* before agent j. The *ex ante* welfare on link l_{ij} can be computed as

$$W_{l_{ij}} = \int_0^\infty e^{-\rho t}(\mu_i^0 + \mu_j^0 - 2c)dt \tag{42}$$

$$-E_{\zeta_i(t^*)}[\int_{t^*}^\infty e^{-\rho t} E_{q_i,q_j}(q_i + q_j - 2c|t \geq t^*)dt] \tag{43}$$

$$-E_{\zeta_j(t^*)}[\int_{t^*}^\infty e^{-\rho t} E_{q_i,q_j}(q_i + q_j - 2c|t \geq t^*)dt] \tag{44}$$

In the case of $\zeta_i(t^*)$, for any $t \geq t^*$, since the learning has stopped, $E_{q_i,q_j}(q_i - c|t \geq t^*) = c - \delta - c = -\delta$ by the definition of t^*. Similarly, $E_{q_i,q_j}(q_j - c|t \geq t^*) > c - \delta - c = -\delta$. Let $h(\delta, t^*) = E_{q_i,q_j}(q_i + q_j - 2c|t \geq t^*) = E_{q_i,q_j}(q_j - c|t \geq t^*) - c - \delta$. We will show that for any t^*, $h(\delta, t^*) < 0$ when δ is sufficiently large. To prove this, we will show that $E_{q_i,q_j}(q_i + q_j - 2c|t \geq t^*)$ is bounded above for any t^*. Consider any *ex post* realization conditional on that it does not hit $c - \delta$ before t^*, there are two possibilities,

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

Clearly, $E(q_j|\zeta_1) > E(q_j|\zeta_2) = c - \delta$. Hence $E_{q_i,q_j}(q_i + q_j - 2c|t \geq t^*) < E(q_j|\zeta_1)$. When $\delta \to \infty$, all realizations never hit the boundary, so $\lim_{\delta \to \infty} E(q_j|\zeta_1) = \mu_j^0$. Therefore, $\forall \epsilon > 0$, there exists δ'_{ij} such that $\forall \delta > \delta'_{ij}$, $|E(q_j|\zeta_1) - \mu_j^0| < \epsilon$. Hence, $\forall \epsilon > 0$, we let $\bar{\delta}_{ij} = \max\{\delta'_{ij}, \mu_j^0 - c + \epsilon\}$, and we ensure for all $\delta > \bar{\delta}_{ij}$, $E(q_j|\zeta_1) - c - \delta < 0$ which also implies that $h(\delta, t^*) < 0$ for all t^*. By choosing $\bar{\delta} = \max_{i,j} \bar{\delta}_{ij}$, we ensure the social welfare is also greater than W^*.

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

divide $H(\delta)$ into two parts,

$$H(\delta) = E_{t_i^* < \hat{t}(\delta)} \left[\int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c | t \geq t_i^*)] dt \right] \tag{45}$$

$$+ E_{t_i^* \geq \hat{t}(\delta)} \left[\int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c | t \geq t_i^*)] dt \right] \tag{46}$$

$$= H'(\delta) + H''(\delta) \tag{47}$$

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

For a given δ_n, there exists $t'(\delta_n)$ such that $\forall t \geq t'(\delta_n)$, $P(t^* < t') < \frac{1}{\delta_n^2}$. Let $\hat{t}(\delta_n) = \max\{t'(\delta_n), \delta_n\}$. Consider $H'(\delta_n)$, it is bounded by

$$|H'(\delta_n)| < P(t_i^* < \hat{t}) \sup_{t_i^* < \hat{t}} | \int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c | t \geq t_i^*)] dt| \tag{48}$$

$$< \frac{\sup_{t_i^* < \hat{t}} | \int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c | t \geq t_i^*)] dt|}{\delta_n^2} \tag{49}$$

$$< \frac{1}{\rho \delta_n^2} \sup_{t_i^* < \hat{t}} |E q_j - c - \delta_n | t \geq t_i^*| \tag{50}$$

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

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

$$|H''(\delta_n)| < P(t_i^* \geq \hat{t}) \sup_{t_i^* \geq \hat{t}} | \int_{t_i^*}^{\infty} E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c | t \geq t_i^*)] dt| \tag{51}$$

$$< \frac{1}{\rho e^{\rho t'}} \sup_{t_i^* > \hat{t}} |E_{q_i,q_j}[e^{-\rho t}(q_i + q_j - 2c | t \geq t_i^*)]| \tag{52}$$

$$< \frac{1}{\rho e^{\rho \delta_n}} \sup_{t_i^* < \hat{t}} |E q_j - c - \delta_n | t \geq t_i^*| \tag{53}$$

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

REFERENCES

[1] D. J. Watts and S. H. Strogatz, “Collective dynamics of ’small-world’ networks,” *nature*, vol. 393, no. 6684, pp. 440–442, 1998.

[2] J. K. Goeree, A. Riedl, and A. Ule, “In search of stars: Network formation among heterogeneous agents,” *Games and Economic Behavior*, vol. 67, no. 2, pp. 445–466, 2009.

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

[4] A. Galeotti, S. Goyal, and J. Kamphorst, “Network formation with heterogeneous players,” *Games and Economic Behavior*, vol. 54, no. 2, pp. 353–372, 2006.

[5] V. Bala and S. Goyal, “A noncooperative model of network formation,” *Econometrica*, vol. 68, no. 5, pp. 1181–1229, 2000.

[6] A. Galeotti, S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv, “Network games,” *The review of economic studies*, vol. 77, no. 1, pp. 218–244, 2010.

[7] M. O. Jackson and A. Wolinsky, “A strategic model of social and economic networks,” *Journal of economic theory*, vol. 71, no. 1, pp. 44–74, 1996.

[8] A. Watts, “A dynamic model of network formation,” *Games and Economic Behavior*, vol. 34, no. 2, pp. 331–341, 2001.

[9] L. Wang and K. Pötzelberger, “Boundary crossing probability for brownian motion and general boundaries,” *Journal of Applied Probability*, pp. 54–65, 1997.

[10] M. Van Der Schaar and S. Zhang, “A dynamic model of certification and reputation,” *Economic Theory*, pp. 1–33, 2014.