NETWORK EVOLUTION WITH INCOMPLETE INFORMATION AND LEARNING

Jie Xu[*], *Simpson Zhang*[†], *Mihaela van der Schaar*[*]

[*]Department of Electrical Engineering, [†]Department of Economics, UCLA

ABSTRACT

We analyze networks that feature reputational learning: how links are initially formed by agents under incomplete information, how agents learn about their neighbors through these links, and how links may ultimately become broken. We show that the type of information agents have access to, and the speed at which agents learn about each other, can have tremendous repercussions for the network evolution and the overall network social welfare. Specifically, faster learning can often be harmful for networks as a whole if agents are myopic, because agents fail to fully internalize the benefits of experimentation and break off links too quickly. As a result, preventing two agents from linking with each other can be socially beneficial, even if the two agents are initially believed to be of high quality. This is due to the fact that having fewer connections slows the rate of learning about these agents, which can be socially beneficial. Another method of solving the informational problem is to impose costs for breaking links, in order to incentivize agents to experiment more carefully.

1. INTRODUCTION

Optimally locating agents within networks is an important problem for many applications, such as organizational structuring, social networks, academic networks, etc [1]. Since agents have costs associated with maintaining links, it is optimal for agents to be connected only if the agents are of sufficiently high quality [2]. But the problem of designing optimal networks is compounded by incomplete information, as the true qualities of agents are often not known [3]. For instance, there is usually only incomplete initial information about the true value of a worker in a company or of a friend in a social network. Even if the initial belief about an agent's quality is favorable, it could turn out that the agent is actually of low quality after more information is learned.

Learning about quality will naturally occur over time as agents form links with each other and send information [3]. As a worker collaborates with his peers and produces some output for example, other workers and the company can update their assessments of his quality. Thus when designing a network, it is necessary to consider both the initial beliefs about agent quality as well as how agents will learn about each other once the network is formed. Such learning can have a strong impact on the resulting network evolution [4][5]. In our model, we study scenarios in which agents learn about each other and break off links with agents with bad qualities, but are unable to create links with good agents due to physical constraints.

Likewise, the shape of the network also strongly affects the rate of learning [6]. Agents who have many links will send more information than agents who have few links, because they are producing more signals. It is easier to learn about the quality of a person who interacts with many other people than a person that interacts with few. And if an agent develops a bad reputation and has all its neighbors sever the links to itself, then learning about this agent can stop completely. Such agents are ostracized from the community as no one is willing to connect with them, and as such these agents have no opportunity to send further information to improve their reputations. Thus the evolution of a network and the learning within that network are intricately intertwined, with one having a direct impact on the other.

Given the close nature of network evolution and learning, it is important to consider the implications that learning has on welfare, both at the individual agent level as well as for the overall network. In our model, we can compute welfare explicitly for any initial network. We show that learning can actually have a negative impact on social welfare in a wide variety of cases, the reason being that faster learning causes agents to disconnect with their neighbors faster if their neighbors send bad information. While this is good for the agent itself, this is bad for the neighbor as the neighbor would now get ostracized from the community and be unable to reap the benefits of the network. On aggregate, this can also be bad for network welfare as a whole: if everyone is learning faster about others, then everyone has the potential to be kicked out of the network sooner themselves as well.

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 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. Thus, it is important to carefully consider which agents are connected with which other agents when designing a network and only considering the connectivity degree is not sufficient. In some cases, a star or a core-periphery network would generate higher social welfare than a complete network [7][8][9].

A potential method of addressing the negative effect of learning is to impose costs on the agents for breaking links with others. For example, this could be introduced in the form of a social stigma, or a deliberate company punishment. By imposing costs on the agents, agents will only cut off links with neighbors once their neighbor's reputation falls extremely low, not just to a moderately low level. Now if a link gets broken and an agent becomes ostracized, it will only be because that agent's contribution to the network is indeed very low, and so the social welfare of the network as a whole can improve as learning becomes faster.

2. RELATED WORKS

Existing works in the network formation literature have studied what networks are formed under complete information when agents know each other's qualities. For example, agents of homogeneous qualities were studied in [10][2][3] and agents of heterogeneous qualities were studied in [7][8][9]. In these models, agents are aware of all the payoff parameters of the network and there is no learning. However, the complete information assumption rarely holds in real world. Few attempts [11] have been made towards studying network formation under incomplete information. Prior work [11] proposes a simplistic model and imposes strong assumptions. For instance, once two agents meet, one agent knows the exact quality of the other agent and also the exact qualities of all indirectly linked agents including those whom it has not met before. Moreover, no rigorous social welfare result is derived due to the intractability of the model. In contrast, the tractability of our proposed model allows us to explicitly compute the social welfare of different network structures even with incomplete information, which lets us compare the optimality of different network structures.

This work is also related to a different strand of literature on social learning [12][13][14][15] and consensus/gossip/diffusion algorithms [16][17][18][19]. Networks in these works are exogenously determined. In contrast, networks in our work are endogenously determined over time as agents are learning and form/sever links. Moreover, the goal of learning in the literature is to learn an exogenous variable such as the underlying environment state while the goal of learning in our work is to learn neighboring agents' qualities to decide whether form or sever links with them.

3. SYSTEM MODEL

We consider an infinite horizon continuous time model with a finite number of agents. Let $V = \{1, 2, ..., N\}$ denote the set of agents. Agents are initially connected according to a network $G = \{g_{ij}\}_{i,j \in V}$ where $g_{ij} = 1$ if agent i and j are connected with each other and $g_{ij} = 0$ otherwise. This initial network G may represent physical constraints and may also be planned, e.g. by the human resource department in a company. Agents that are not linked initially can never become linked. We say that agent i and j are neighbors if $g_{ij} = 1$. At each moment in time $t \geq 0$, each agent i can break any of its links unilaterally. Hence, the network will evolve over time. Let G^t be the network at time t and $k_i^t = \sum_j g_{ij}^t$ be the number of links that agent i has at time t. Note that $G^0 = G$. Linking with other agents is costly. An agent i must pay a flow cost c for each of its links, which represents the agent's cost of maintaining the link. Hence, at time t, the cost that agent i must pay is $k_i^t c$.

Agents also obtain benefits from their links, depending on the neighbors' qualities. Each agent i in this network has a fixed quality q_i that is determined at the start of the game according to a commonly known normal distribution $Norm(\mu_i, \sigma_i^2)$ with $\mu_i > c$. This quality q_i is not known to all agents and we do not require that agent i knows its own quality either. When agent i is linked, it generates (random) flow benefits b_i^t for all agents that are linked with it at time t and these benefits are (locally) publicly observable[1]. This signal of flow benefits thus represents the information sent by agent i from which other agents can learn i's true quality. These benefits are noisy, and the evolution of these benefits follows a Brownian motion $dB(t) = q_i dt + v_i(k_i^t) dZ(t)$ where the drift rate is the true quality q_i, the instantaneous volatility rate $v_i(k_i^t)$ 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. In particular, the volatility rate $v_i(k_i^t) = (k_i^t \tau_i)^{-1/2}$ where τ_i is the base precision of the Brownian motion of agent i's benefit process. The more links an agent has, the lower the volatility rate of its benefit Brownian motion and hence, learning about its true quality q_i will be faster. This captures the fact that agents who have many links will send more information than agents who have few links.

Agents are myopic and will maintain the link if and only if they believe that the current benefit of linking with another agent exceeds the link maintenance cost. Since the flow benefits are (locally) publicly observable, neighbors of agent i will have a common belief of any agent i's quality. We define the agent's benefit history as the history of all previous benefits, $\mathcal{H}_i^t = \{b_i^{t'}\}_{t'=0}^t$. If at a time t all links of agent i are severed, then no benefit will be produced by agent i and this will be denoted as $b_i^t = \emptyset$. Note that in this case no information is added and hence, the Brownian motion of agent i is stopped at the current level. As mentioned, agents have a prior belief of an agent i's quality $Norm(\mu_i, \sigma_i^2)$ and update this belief in a Bayesian fashion in light of the observations of flow benefits. This signal 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

[1] We only require these benefits be observed by direct neighbors.

[20]. An agent j will choose to maintain the link with agent i if and only if $E[q_i|\mathcal{H}_i^t] > c$ with the expectation taken over the posterior belief distribution. We denote $\mu_i^t \triangleq E[q_i|\mathcal{H}_i^t]$ and call it the *reputation* of agent i at time t. It can be shown [21] that μ_i^t is a martingale given the information available to agents and hence the belief is correct.

In this model, agents' true qualities are unknown a priori and their generated flow benefits represent incomplete information from which other agents can learn their true qualities. This learning leads to network evolution in which links are broken over time and the flow benefits thus change, affecting both individual agent welfare as well as the overall network welfare. In the next sections, we will study the table networks, the implications that learning has on welfare, and how the shape of the initial network affects the rate of learning and the welfare.

4. STABLE NETWORKS

The network G^t is evolving over time due to agents breaking links with each other. We call the limiting network structure when t tends to infinity, denoted by G^∞, a *stable* network. Because all neighbors of an agent i break the link with i at the same time due to the common belief, once the links of i are broken, agent i will be ostracized from the network forever in the future since its reputation will stay at the current level and will never go up. Therefore, the network is always shrinking, i.e. $G^t \supseteq G^{t'}, \forall t < t'$. However, G^∞ is not always an empty network as we will show shortly.

To understand what G^∞ can be, we need to investigate whether a link $l_{ij}, \forall i, j$ still exists at $t = \infty$. A link l_{ij} still exists if and only if neither agent i's or j's reputation hits c before $t = \infty$ (otherwise, the link is severed by either i or j unilaterally). Let ε_i^t denote the event that agent i' reputation hits c at time t. Then the probability that a link l_{ij} exists in a stable network is $P(\varepsilon_i^\infty)P(\varepsilon_j^\infty)$. Using standard mathematical results of Brownian motion hitting probabilities, $P(\varepsilon_i^\infty)$ can be shown to be independent of the initial network structure and can be explicitly calculated [22].

Theorem 1. *A network structure G^∞ can be a stable network if and only if $G^\infty \subseteq G^0$ and $g_{ij}^\infty = \mathbf{1}(\varepsilon_i^\infty)\mathbf{1}(\varepsilon_i^\infty)$. Moreover, if this condition is satisfied, then G^∞ will be a stable network with probability $\prod_i P(\varepsilon_i^\infty)$.*

5. IMPACTS OF LEARNING

5.1. Welfare

Each agent generates flow benefits for its link(s) and at the same time also sends implicit information about its true quality. The speed of sending information and hence the speed of learning is determined by the base precision of an agent's Brownian motion. The higher the precision, the faster the information diffuses. In this section, we study how learning in terms of agents' base precision affects the welfare.

The individual welfare of an agent i is the *ex ante* expected long-term payoff that it obtains from all of its links. Let j be i's initial neighbor, we denote W_{ij} as the *ex ante* expected long-term payoff obtained by agent i from the link with agent j, which can be computed as

$$W_{ij} = \int_{-\infty}^{\infty} \int_{0}^{\infty} e^{-\rho t}(q_i - c)P(l_{ij}^t|\boldsymbol{q}, G)dt\phi(\boldsymbol{q})d\boldsymbol{q} \quad (1)$$

where $\phi(\boldsymbol{q})$ is the joint normal distribution of all agents' quality $\boldsymbol{q}$, $P(l_{ij}^t|\boldsymbol{q}, G)$ is the survival probability of the link between i and j and $\rho > 0$ is the discount rate. Hence, agent i's welfare is $W_i = \sum_{j:g_{ij}=1} W_{ij}$.

The next theorem shows how agent i's base precision τ_i affects its own welfare.

Theorem 2. *For any initial network G, an agent i's welfare is decreasing in its own base precision τ_i.*

Proof Sketch. Consider any *ex post* realization of agents' reputation hitting time events $\varepsilon = (\varepsilon_1^{t_1}, ..., \varepsilon_N^{t_N})$. The long-term payoff of agent i depends on the hitting time t_i of itself and its neighbors whose reputations never hit c before infinity. For neighbors whose reputations do hit c before infinity, the expected value of linking with them would be zeros, so they can be ignored. Increasing agent i's own precision decreases its own hitting time but does not change the set of neighbors whose reputations never hit c. This weakly decreases the long-term payoff of agent i in any *ex post* realization, with the decrease being strict if the event $\epsilon_i^{t_i}$ with $t_i < \infty$ occurs. Therefore, agent i's welfare is strictly decreasing in its own precision τ_i. $\square$

Theorem 2 proves that increasing one's precision (hence the information diffusion speed) is harmful to its own welfare. This is because, in any *ex post* realization, the hitting time of its reputation against c becomes sooner and hence the agent is ostracized from the network sooner.

An agent's information sending speed also affects other agents' welfare in the network. Faster learning about agent i's quality causes agent i to disconnect with its neighbors sooner. However, whether learning about i leads to greater or lower welfare for i's neighbors depends on the specific network. Theorem 3 shows that for any network G without cycles, any agent's information sending speed has a positive impact on its neighbors' welfare.

Theorem 3. *For any initial network without cycles, increasing any agent i's base precision τ_i increases its neighbor j's welfare.*

Proof Sketch. Suppose i and j are any two neighbors. Since there are no cycles in the initial network, there is no other path connecting i and j except their direct link. Hence, G can be partitioned into two subnetworks G_i and G_j where $i \in G_i$, $j \in G_j$ and there is no overlap between G_i and G_j

(expect that i and j are linked). Consider any *ex post* realization $\varepsilon = (\varepsilon_1^{t_1}, ..., \varepsilon_N^{t_N})$. In the case of $t_j = \infty$, changing agent i' precision τ_i does not affect agent j' hitting time and hence has no impact on the hitting times of agents in G_j. Therefore, agent j's payoff is not affected. In the case of $t_j < \infty$, increasing τ_i decreases agent i's hitting time which results in one link of agent j being severed sooner. This leads a slower learning of agent j's quality and hence its hitting time becomes later. Therefore, agent j's payoff becomes larger. □

Networks without cycles, such as stars and trees, are common networks in practice. In such networks, the result of losing a link with a neighbor is to reduce one of its neighbor's links and thereby reduce the rate at which the neighbor sends information. This in turn increases the welfare of that neighbor. However, when cycles are present in a network, information sending and link breaking has a more complicated impact on the learning rates of other agents' qualities. Thus the result in Theorem 3 may not hold. A counter example is a network with many relay agents with very low initial expected qualities as shown in Figure 1. Suppose agent 3's initial expected quality μ_3 is very high and the initial expected qualities of agent 4 to K are very close to c. In this network, increasing agent 1's precision makes its reputation hitting time sooner and thus slows down the learning about agent 4 to K's qualities. Therefore, agent 2 has many links for a longer time which makes the learning about its quality faster overall. Hence, agent 2's long-term payoff becomes smaller with the increase of agent 1's precision.

Theorem 4. *For initial networks with cycles, it is possible that increasing some agent i's base precision τ_i decreases its neighbor j's welfare.*

Proof. An informal proof uses the above example. □

Now we study how learning affects the network welfare, which is defined as the sum of individual welfare $W = \sum_i W_j$.

Theorem 5. *Given any network G, multiplying all agents' base precisions by the same number $d > 1$ decreases the network welfare.*

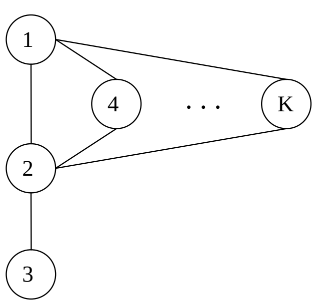

Fig. 1. An illustrative example of networks with cycles.

Proof. This is a direct result of Theorem 2. □

Theorem 5 is due to the fact that if everyone is learning faster about others, then everyone has the potential to be kicked out of the network sooner themselves as well, which decreases the network welfare. Thus the theorem implies that learning actually has a negative impact on network welfare.

5.2. Implications for Network Planning

As we have shown, the shape of the network affects the rate of learning and hence the achievable social welfare, so it is natural to study which initial network structure yields the highest social welfare. This is an important topic for organizational structuring, social networks, academic networks etc., when the exact qualities of agents are unknown but only the distributions of the qualities are known. Here we show that a complete network in which all agents are connected with each other is not necessarily the optimal structure in the presence of incomplete information and learning, no matter how many agents there are.

Theorem 6. *For any number N of agents, a complete initial network may not be the optimal initial network.*

Proof Sketch. We prove by constructing a counter example. Suppose that the initial expected quality of agent 1 μ_1 is very large and the initial expected qualities of other agents $\mu_i \approx c, \forall i \neq 1$. Hence, linking with agent 1 produces high payoffs while linking with the remaining agents produces negligible payoff. Consider any *ex post* realization of agents' reputation $\varepsilon = (\varepsilon_1^{t_1}, ..., \varepsilon_N^{t_N})$ in a complete network. Let $\varepsilon' = (\varepsilon_1^{t'_1}, ..., \varepsilon_N^{t'_N})$ be the corresponding realization in a star network with agent 1 being at the center. It can be shown that $t'_1 \leq t_1$ and $t'_i \geq t_i, \forall i \neq 1$. Since the initial mean quality of agent $i \neq 1$ is low, having a shorter t'_1 does not decrease agent 1's welfare much. However, the remaining agents' welfare can be significantly improved because of the increase in t'_i and the fact that agent 1's quality is very high. Therefore the star network can yield a higher social welfare than the complete network. □

Under the proposed model, many other implications on the network planning can be derived. For example, the core-periphery network in which agents with high initial expected qualities compose the core and agents with low initial expected qualities compose the periphery can be better than the complete network in some cases.

5.3. Cost for Breaking Links

If the agents do not observe any information about each other, then agents do not learn each other's true quality and hence all initial links would remain connected forever. Therefore the social welfare without learning can be computed as

$\sum_{l_{ij}} \int_0^\infty e^{-\rho t}(\mu_i + \mu_j - 2c)dt$. We have shown that learning actually has a negative impact on the network welfare in Theorem 5. A potential method of addressing the negative effect of learning is to impose costs on the agents for breaking links with others. For example, this could be introduced in the form of a social stigma, or a deliberate company punishment. Let Δ be such an instantaneous (infinitesimal) cost. In this case, agents will only cut off links with neighbors once their neighbors' reputation falls below $c - \Delta$.

Theorem 7. *For any network, there exists Δ large enough such that the achievable social welfare with learning is greater than that without learning.*

Proof Sketch. The social welfare with learning can be computed by

$$\begin{aligned} W = &\sum_{l_{ij}} \int_0^\infty e^{-\rho t}(\mu_i + \mu_j - 2c)dt \\ &- \sum_{l_{ij}} E_{\zeta(t_i^*)}[\int_{t_i^*}^\infty e^{-\rho t} E_{q_i,q_j}[(q_i + q_j - 2c|t \geq t_i^*)]dt] \\ &- \sum_{l_{ij}} E_{\zeta(t_j^*)}[\int_{t_j^*}^\infty e^{-\rho t} E_{q_i,q_j}[(q_i + q_j - 2c|t \geq t_j^*)]dt] \end{aligned} \tag{2}$$

where $\zeta(t_i^*)$ is the event in which agent i's reputation hits $c - \Delta$ before agent j at time t_i^*. Since the achievable social welfare without learning is $\sum_{l_{ij}} \int_0^\infty e^{-\rho t}(\mu_i + \mu_j - 2c)dt$, in order to make W greater, we need to make the last two terms negative. Fix a t_i^*, then $E(q_i|t \geq t_i^*) = \mu_i^{t_i^*} = c - \Delta$ and $E(q_j|t \geq t_i^*)$ is bounded above. Hence, by choosing Δ large, we can ensure that $E(q_i + q_j - 2c|t \geq t_i^*)$ is negative for all $t_i^* > 0$. Therefore the social welfare with learning is greater than that without learning. □

6. CONCLUSION

This paper studies a framework for network evolution with incomplete information and learning. We showed what the stable networks could be and how learning affects both the individual and the social welfare when agents' qualities are unknown and must be learned through past observations. Some important implications on optimal network planning in the presence of incomplete information and learning are also revealed but many more can be derived using the current model.

7. REFERENCES

[1] Matthew Haag and Roger Lagunoff, "Social norms, local interaction, and neighborhood planning," *International Economic Review*, vol. 47, no. 1, pp. 265–296, 2006.

[2] Venkatesh Bala and Sanjeev Goyal, "A noncooperative model of network formation," *Econometrica*, vol. 68, no. 5, pp. 1181–1229, 2000.

[3] Alison Watts, "A dynamic model of network formation," in *Networks and Groups*, pp. 337–345. Springer, 2003.

[4] Ceyhun Eksin, Pooya Molavi, AR Ribeiro, and Ali Jadbabaie, "Learning in network games with incomplete information: Asymptotic analysis and tractable implementation of rational behavior," *Signal Processing Magazine, IEEE*, vol. 30, no. 3, pp. 30–42, 2013.

[5] Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[6] S Nageeb Ali and David A Miller, "Ostracism," *Unpublished manuscript, University of California at San Diego*, 2012.

[7] Andrea Galeotti and Sanjeev Goyal, "The law of the few," *The American Economic Review*, pp. 1468–1492, 2010.

[8] Yu Zhang and Mihaela van der Schaar, "Information production and link formation in social computing systems," *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 11, pp. 2136–2145, 2012.

[9] Yu Zhang and Mihaela van der Schaar, "Strategic networks: Information dissemination and link formation among self-interested agents," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 6, pp. 1115–1123, 2013.

[10] Matthew O Jackson and Asher Wolinsky, "A strategic model of social and economic networks," *Journal of economic theory*, vol. 71, no. 1, pp. 44–74, 1996.

[11] Yang Song and Mihaela van der Schaar, "Dynamic network formation with incomplete information," in *XVI Southwest Economic Theory Conference UC Irvine*, 2014.

[12] Douglas Gale and Shachar Kariv, "Bayesian learning in social networks," *Games and Economic Behavior*, vol. 45, no. 2, pp. 329–346, 2003.

[13] Benjamin Golub and Matthew O Jackson, "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, pp. 112–149, 2010.

[14] Daron Acemoglu and Asuman Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, no. 1, pp. 3–49, 2011.

[15] Vikram Krishnamurthy, "Quickest detection pomdps with social learning: Interaction of local and global decision makers," *Information Theory, IEEE Transactions on*, vol. 58, no. 8, pp. 5563–5587, 2012.

[16] Alexandros G Dimakis, Soummya Kar, José MF Moura, Michael G Rabbat, and Anna Scaglione, "Gossip algorithms for distributed signal processing," *Proceedings of the IEEE*, vol. 98, no. 11, pp. 1847–1864, 2010.

[17] Angelia Nedic, Asuman Ozdaglar, and Pablo A Parrilo, "Constrained consensus," *LIDS Report*, vol. 2779, 2008.

[18] Jianshu Chen and Ali H Sayed, "Diffusion adaptation strategies for distributed optimization and learning over networks," *Signal Processing, IEEE Transactions on*, vol. 60, no. 8, pp. 4289–4305, 2012.

[19] Pooya Molavi, Ali Jadbabaie, Kamiar Rahnama Rad, and Alireza Tahbaz-Salehi, "Reaching consensus with increasing information," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 7, no. 2, pp. 358–369, 2013.

[20] Herman Chernoff, "Optimal stochastic control," *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 221–252, 1968.

[21] Godfrey Keller, "Brownian motion and normally distributed beliefs," Tech. Rep., Working Paper, University of Oxford December, 2009.

[22] Mihaela Van Der Schaar and Simpson Zhang, "A dynamic model of certification and reputation," in *Proceedings of ACM conference on Economics and Computation (EC)*, 2014.