

# RESEARCH STATEMENT

BENJAMIN GOLUB

My research examines social and economic networks, primarily using economic theory, along with some experimental work. Network models focus on the details of who interacts with whom, and offer new insights on particular applications such as social learning, as well as perspectives on fundamental issues in economic theory. This statement discusses examples of my work in several overlapping topics: (1) local spillovers and externalities, as when students’ effort or firms’ innovation directly affects particular “neighbors”; (2) learning and adaptation, as when a community assesses the payoffs of a new opportunity or technology and its members learn from each other; (3) contagion of failure in financial and production networks and the implications for policies to improve resilience.

## 1. SPILLOVERS AND EXTERNALITIES IN NETWORKS

Many questions in topics ranging from organizational economics to the economics of crime can be studied using games with heterogeneous spillovers in incentives. For instance, consider workers who choose levels of effort and have incentives to coordinate these with their direct collaborators (network neighbors), because their efforts are complementary. Suppose their preferred levels of effort also depend on individual attributes—such as how much they value a project. Consider a Nash equilibrium of such a game, defined by a level of effort for each individual. Some basic questions are: (1) How do equilibrium outcomes depend on the structure of the network and individuals’ own attributes? (2) How does welfare depend on the same parameters? (3) How could a planner intervene to achieve, e.g., an increase in average effort, or an improvement in individuals’ welfare?

Social spillovers have long been considered important for understanding aggregate outcomes such as the crime rate (Glaeser, Sacerdote, and Scheinkman, 1996). An early concern of network theory in economics was modeling these spillovers, with detailed modeling of who interacts with whom, as in Goyal and Moraga-Gonzalez (2001) and Ballester, Calvó-Armengol, and Zenou (2006). These studies defined a network where the link from  $i$  to  $j$  describes how much  $i$ ’s best-response action is affected by  $j$ ’s action. They then used this network to characterize the role of various individuals’ attributes in determining aggregate effort.<sup>1</sup> There turned out to be a very close connection between econometric models of peer effects and models of network spillovers (Bramoullé, Djebbari, and Fortin, 2009; Blume, Brock, Durlauf, and Jayaraman, 2015) and these connections have proved useful for applied problems—see, e.g., Calvó-Armengol, Patacchini, and Zenou (2009), Acemoglu, García-Jimeno, and Robinson (2015), and Banerjee, Chandrasekhar, Duflo, and Jackson (2019b).

A theme emerging from these studies is what we might call *the centrality principle*. In settings where the spillovers between agents are given by strategic complements, those who are more network-central<sup>2</sup> have more influence: their attributes play a relatively large role in determining a group’s average or total effort or output. This has also been a useful guiding principle in thinking about whom to target in order to increase or decrease some level of activity: more “ripple effects” emanate from more central agents, and so targeting them often makes a bigger difference.

However, the centrality principle leaves much unresolved. For example, consider games of strategic substitutes, such as local public goods games in networks studied in Bramoullé, Kranton, and d’Amours

---

*Date:* December 11, 2021.

<sup>1</sup>Related models have been important in the literature on spillovers with incomplete information: see, for example, Angeletos and Pavan (2007), Hellwig and Veldkamp (2009), and Bergemann, Heumann, and Morris (2015).

<sup>2</sup>In the network of strategic spillovers mentioned above, often according to classical measures such as eigenvector and Bonacich centrality.

(2014). In such games, increasing the effort of a highly central agent causes many others to want to *decrease* their efforts. So the net effect on total effort could be negative, and the effect on welfare is ambiguous. The received theory offers much less guidance on targeting for behavior change or welfare improvement in such settings.

Galeotti, Golub, and Goyal (2020) show that as we consider natural intervention problems that include, for instance, welfare-optimal intervention with strategic complements, the structure of optimal interventions depends on different and more subtle aspects of a network game. The key method of the paper is as follows: To diagnose which interventions are most effective for increasing a given objective function, we identify how a given profile of incentive changes is amplified or attenuated by the strategic spillovers in the network, and how the planner should take this into account.

Reassuringly, we recover a version of centrality principle in a special case: If actions are strategic complements, the optimal intervention changes all agents' incentives in the same direction and does so in proportion to their eigenvector centralities, at least when the intervention is large enough. Next, we consider games of strategic substitutes—e.g., ones in which an individual has incentives to free-ride on the effort of a neighbor. Here the optimal intervention is very different: it moves neighbors' incentives in *opposite* directions, dividing local communities into positively and negatively targeted agents.

To derive these results and characterize optimal interventions more generally, we introduce a new method of decomposing any potential intervention into orthogonal *principal components* determined by the network. We give a complete description, in terms of these principal components, of how the planner optimally focuses interventions. Welfare-optimal interventions whose structure is rather complex in the original description of the game have a simpler description in the basis given by the principal components. For example, the two different targeting schemes mentioned in the previous paragraph simply correspond to two particular principal components which are “extremes” in a precise sense. By developing a new connection between network games and principal component analysis, we enable the use of new tools from applied mathematics, where the relationship between principal components and the underlying network has been extensively studied.

**Incomplete information.** Several recent papers, such as de Martí and Zenou (2015) and Calvó-Armengol et al. (2015), have pointed out that the presence of incomplete information drastically complicates the analysis of network games. In Golub and Morris (2020a,b), we examine how the analysis of network games can be extended to allow for flexible incomplete information—and how many key insights, suitably adapted, can be extended to this richer domain. For a motivating example, consider again the problem of a collaborating team in an organization, but now with asymmetric information about the returns of the project, and possibly heterogeneous prior beliefs about those returns.

When both asymmetric information and heterogeneous priors are permitted, predicting even average behavior becomes quite subtle. Golub and Morris (2020a) develops an analysis of how coordination motives and asymmetric information jointly determine group behavior when coordination concerns are important. This has some unexpected consequences. For example, suppose a planner would like higher actions to be taken (for instance, because equilibrium actions are inefficiently low). Moreover, suppose there is a member of the group who is very (ex ante) optimistic about the project. One may think, based on the centrality principle mentioned above, that giving the optimist a very central position in the network will maximize his effective influence on others. But in fact, a better course of action can be to leave his centrality unchanged, and instead to *reduce* the precision of his private information. This makes the optimist's behavior stick closer to his optimistic prior; the result is to increase the average action in the coordination game equilibrium, as others best-respond to his optimism. Thus, information can be more important than centrality: making an individual less privately informed, and therefore more “resolute” or “committed” to his ex ante preferred action, can be the best way of making him influential.

Methodologically, the key to our general analysis is the construction of a certain constructed network that encodes *both* the network structure capturing coordination concerns and the structure of incomplete information. We then use network methods to analyze equilibrium behavior. An important theme is that iterated expectations (e.g., A's expectation of B's expectation of C's action) play a crucial role in the

coordination game, and our network formalism allows us to study these using techniques developed for Markov chains. This connection builds on ideas of Samet (1998). The companion paper Golub and Morris (2020b) explores some related methodological issues concerning the behavior of iterated expectations. While Samet (1998) characterized these in the case of common priors, we are able to characterize how iterated expectations behave in a larger class of environments with heterogeneous priors.

**Externalities rather than spillovers: A network perspective on negotiations.** In all the papers discussed in this section, the solution concept has been static Nash equilibrium: all agents choose their actions best responding to others. For instance, in a team game, agents unilaterally decide how much effort to contribute. As mentioned above, in this case what matters for the comparative statics of outcomes is the network of strategic spillovers (complements and substitutes) in equilibrium: how much each agent’s effort affects others’ incentives.

However, in many settings of interest, agents are not playing a noncooperative voluntary-provision game. They may, instead, be involved in a negotiation where they can commit to actions. From a normative perspective, we may simply be interested in when a group can achieve Pareto improvements on a status quo. Then what matters is not strategic interdependencies in a voluntary-provision game (how much  $i$ ’s action affects  $j$ ’s noncooperative best response), but the structure of externalities: how much  $i$ ’s action affects  $j$ ’s welfare.

Elliott and Golub (2019) described how a network perspective on the externalities sheds light on efficient solutions. More concretely, suppose that agents (e.g., countries) can pollute less at a net private cost. Due to geography and other asymmetries, the benefits of such an effort are not distributed uniformly. We show how a planner can use a network that describes the externalities to analyze Pareto-improvements, and efficient negotiated outcomes, available to the group.

To describe the essence of this theory, we consider a simple example. Suppose agent  $A$  can, by making a sacrifice, make  $B$  better off, but  $B$ ’s concessions create no direct value for  $A$ . Even in repeated interaction, these two cannot improve on the status quo by favor-trading without recourse to some other currency. But if  $B$  can help  $C$ , who can in turn help  $A$ , then there is scope for Pareto-improving cooperation among the three, which can be interpreted as committing to favor trading: each is willing to do a favor for its beneficiary to sustain the cycle and get a favor (from someone else). Thus, seeking sufficiently strong cycles of benefit flows (strong enough to cover everyone’s costs of contributing) is key to structuring Pareto improvements. Motivated by this idea, we develop the above observation about cycles into a general characterization of whether and how much a group can improve on the status quo. The key is to analyze a matrix whose entries record the marginal benefits per unit of marginal cost that each agent can confer on each other, for a given action profile. A certain network statistic, the *spectral radius* of the externality network quantifies the collective returns to increasing actions. We use this measure to characterize players who are essential to negotiations. We also describe when negotiations can be subdivided without much loss. A negotiation can be efficiently subdivided when the marginal externalities network is nearly disconnected. This intuitive observation turns out to be more subtle than it seems; more importantly, finding good subdivisions can be reduced to well-studied problems in network partitioning (Spielman, 2007).

Beyond studying whether cooperation can be sustained, we also focus on certain negotiated outcomes that we motivate with bargaining models (Yildiz, 2003; Penta, 2011). We show that these outcomes—the Lindahl outcomes—are exactly characterized by a centrality property: agents contribute in proportion to their centrality in the externality network that reflects who receives benefits from others’ contributions. Thus, in a precise sense (and without reliance on parametric assumptions) we show that a certain important class of negotiated outcomes can be characterized in terms of centralities in an underlying externality network.

A conceptual contribution of this paper is that certain network statistics that are familiar from network games analysis appear in the analysis of Pareto efficient outcomes. However, they play very different roles from the familiar ones in network game theory. This paper thus opens up the potential for network theory to offer new tools, both conceptual and computational, for the design of negotiations or favor-trading markets.

## 2. SOCIAL LEARNING AND INFORMATION AGGREGATION IN NETWORKS

People learn about unknown states (e.g., the quality of a new technology) by observing others' decisions as well as by directly sharing opinions.<sup>3</sup>

These questions have motivated several programs of study. One thread, on Bayesian models, includes the sequential observational learning papers of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992); the work of Vives (1993) on the rate of social learning in a rational expectations equilibrium; and large literatures that developed on related topics.<sup>4</sup> A parallel literature developed around the observation that rule-of-thumb, behavioral learning rules may be more realistic in complex environments. Ellison and Fudenberg (1993, 1995) examined different naive rules of thumb for social learning and argued that these could be reasonably efficient; Bala and Goyal (1998) introduced a boundedly rational model of learning in a network. There is a vibrant literature on these subjects, including in behavioral economics—see, e.g., Eyster and Rabin (2010, 2014) and the recent survey of Levy and Razin (2019). A question that recurs throughout much of the work in this literature is the extent to which learning is efficient or aggregates information well, leading to estimates that closely approximate the truth—and how this depends both on the structure of information and observation.

Golub and Jackson (2010) was a paper in the rule-of-thumb paradigm, focusing on how *network structure* matters for the quality of information aggregation. Substantively, it found that some network structures allow naive agents to achieve very good aggregation in large populations, while others do not, and gave some general characterizations of when each case obtains. Methodologically, a key novelty of the paper was that it showed that this question is closely related to certain properties of network centrality distributions in large networks—in particular, the distribution of eigenvector centrality, a canonical statistic in network theory. A substantial body of work has subsequently built on this observation, bringing together learning theory and network theory; below I briefly discuss this literature.

In the model, there is a real-valued uncertain state (say, the best action to take) that is drawn once and for all. A population of  $n$  agents in a network receive private signals about the state and then repeatedly update their estimates of it, based on the estimates of their network neighbors. The updating occurs according to the simple weighted-averaging rule of DeGroot (1974): agents form their opinions by averaging their previous opinions with those of their contacts. This can be motivated as a natural form of “persuasion bias” in which people do not optimally account for the fact that their neighbors have correlated opinions conditional on the state because they share common influencers (see DeMarzo, Vayanos, and Zwiebel (2003) and Eyster and Rabin (2010) for detailed discussions of this interpretation). We focus on the long-run opinions: what agents believe after many rounds of such updating.

We ask when these long-run estimates are close to the true state in a large network. The first result is that a sequence of networks<sup>5</sup> achieves good aggregation if and only if the maximum eigenvector centrality of any agent tends to 0. The second set of results gives various more primitive conditions for this decaying-centrality condition to hold. A fundamental obstruction is the existence of finite sets of agents that are “prominent”—influencing many others directly or indirectly.<sup>6</sup> This paper also made extensive use of the fact that a node's influence on group opinion in the DeGroot model is equal to its eigenvector centrality, and contributed to diffusing this observation in the networks literature.

---

<sup>3</sup>The resulting information flows play a role in a variety of economically important processes, such as job search (Calvo-Armengol and Jackson, 2007; Beaman and Magruder, 2012), household financial planning (Dufo and Saez, 2003), and the choice of technologies or production methods (BenYishay and Mobarak, 2019).

<sup>4</sup>To mention just a few other contributions in the Bayesian literature especially relevant to networks, Banerjee and Fudenberg (2004) moved away from the “single file” paradigm of standard sequential models, while Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) modeled networks explicitly.

<sup>5</sup>We consider such sequences so that we can make sense of the statement “in a large network.”

<sup>6</sup>We also give several more intricate sets conditions. The question of when a sequence of networks has good aggregation is equivalent to the question of when a sequence of finite Markov chains (of increasing size) has stationary distributions that converge pointwise to the zero vector; I believe it is still not known exactly what this means in terms of the graph structure of the Markov chains.

Golub and Jackson (2012b) continued to study rule-of-thumb learning in large networks, but rather than focusing on the *quality* of aggregation, this paper examined the *speed* of convergence to a consensus estimate. The main finding of the paper is that a key determinant of speed is homophily, the tendency of people to socialize most with demographically similar others. Formally, the speed of convergence is characterized in terms of a group-level measure of segregation called *spectral homophily*. Moreover, disagreement takes the form of polarization of opinions across groups: societies that exhibit network segregation along demographic lines converge internally, but persistently disagree with each other. There are two important steps: (1) to describe the dynamics of updating using the matrix algebra of Markov chains (e.g., Levin, Peres, and Wilmer, 2009) and give a geometric description of the “main component” of disagreement in opinion space, following DeMarzo, Vayanos, and Zwiebel (2003); (2) use the spectral properties of random graphs (e.g., Chung, Lu, and Vu, 2004) to express these rates in terms of intuitive statistics of social structure. Thus, methodologically, this study strengthened connections between a central concept in sociology and two active applied mathematics literatures, while also developing new, empirically relevant network measures. The characterization just discussed is reliant on suitable statistical assumptions about the underlying network. In Golub and Jackson (2012a) we found that these statistical assumptions seem to hold up in the data on high school friendship networks from the National Longitudinal Study of Adolescent to Adult Health. Jackson (2021) surveys how homophily measures and consequences of the type explored in our work have percolated into policy discussions of inequality and social segregation.

Work on behavioral learning models in social networks has developed actively since these papers were circulated. I briefly highlight several threads that are most closely connected to the concerns of the above papers. One strand (e.g., Molavi, Tahbaz-Salehi, and Jadbabaie, 2018) has examined generalizations of the DeGroot learning rule, characterized classes of rules in terms of their essential behavioral predictions (i.e., axiomatized them), and analyzed the properties needed for good or fast aggregation of information. The aggregation properties of the DeGroot rule established in Golub and Jackson (2010) have served as a benchmark, but the qualitative requirements for good learning under behavioral rules are now much better understood.<sup>7</sup> Other work (e.g., Tahbaz-Salehi and Jadbabaie, 2009; Acemoglu, Ozdaglar, and ParandehGheibi, 2010; Banerjee, Chandrasekhar, Duflo, and Jackson, 2019b) has considered richer (sometimes stochastic) processes of meeting or updating and examined the extent to which the eigenvector-centrality characterization of influence can be used in other contexts. Golub and Sadler (2016) surveyed progress in this area, covering work in the Bayesian paradigm as well.

Of course, it is also important to understand how people actually update their estimates and learn from each other. There has been an active experimental literature comparing rational and behavioral models in networks, and asking whether real people are able to counteract the influence of overly central agents—see, e.g., Mobius, Phan, and Szeidl (2015) and Chandrasekhar, Larreguy, and Xandri (2020).

I close this section by discussing my most recent theoretical paper on learning in networks. The standard economic models in social learning have tended to take the “state of the world” as fixed and examine long-run properties of opinion dynamics against the backdrop of this fixed state. Realistically, the state of the world—the right technology to use, conditions in a market, etc.—is often changing at a rate comparable to the rate at which social learning changes agents’ estimates. Such a situation requires a different model. In the model of Dasaratha, Golub, and Hak (2020), the state of the world changes over time as agents constantly receive signals about it, with each agent having a potentially a different distribution of private signals, which can be thought of as a “perspective.” Agents form estimates using their own signals and the recent estimates of their peers; we study an equilibrium in which all agents are optimizing against the (endogenous) peer learning rules.<sup>8</sup> We identify informational environments in which simple linear updating rules reminiscent of DeGroot’s heuristic are actually part of an equilibrium. The main substantive finding is that information aggregation is good as long as each individual has access to a set of neighbors that is sufficiently diverse, in the sense of having enough different perspectives represented in substantial numbers among these neighbors. If individuals’ neighborhoods are *not* diverse, then social learning is

<sup>7</sup>Mossel, Sly, and Tamuz (2015) characterizes asymptotic *rational* learning in large networks and some of the key conditions turn out to be related to ones we studied in Golub and Jackson (2010)!

<sup>8</sup>For an engineering perspective on a changing-state learning problem, see Shahrampour, Rakhlin, and Jadbabaie (2013).

inefficiently confounded and far from optimal. The essential intuition is that a lack of signal diversity creates identification problems. If one’s neighbors have similar signals, then they use information similarly; that situation (reminiscent of collinearity in statistics) makes it harder to figure out the new developments behind their behavior. Diversity is, in a sense, more important than precision: giving everyone better signals can hurt aggregation severely if it makes those signals homogeneous.

**Endogenous engagement in social learning: An experimental agenda.** An obvious question about social learning is how it is affected by endogenous acquisition of information. Though social learning, as mentioned above, can be very important for high-stakes decisions, asking questions is often essential to get social learning started. Beyond the fact that social learning may involve physical or information-processing costs—an issue studied by Galeotti and Goyal (2010), Niehaus (2011), and Acemoglu, Bimpikis, and Ozdaglar (2014), among others—seeking information may have a different kind of cost, associated with others’ perceptions (or, in other words, with one’s social image). Being seen not to know something can in itself be compromising. An experimental agenda I am working on focuses on the distinctive impediment to social learning that arises from image concerns. It examines how, in light of such concerns, agents decide whether to participate in information exchange, whom to talk to, and what to ask.

Chandrasekhar, Golub, and Yang (2019) posits a signaling model in which people are instrumentally concerned about their reputation for ability. They refrain from asking questions to avoid updating others’ beliefs in a bad direction. Identifying instrumental ability-signaling can be challenging because of a distinct but observationally similar effect that we call *shame*: the reluctance to interact with people who have bad beliefs about you, no matter how they got those beliefs. In a simple model, we distinguish the observable effects of signaling and shame and describe how they depend (differently) on the environment. To investigate whether and how much they can inhibit learning in practice, we run a field experiment with over 1200 subjects. We find that, combined, image concerns can severely deter seeking behavior. We show that the shame effect is particularly pronounced among socially close individuals (in terms of network distance and caste co-membership), whereas the signaling effect dominates among pairs who are less close.

The next step is to examine the implications for how policymakers should disseminate information. Should they broadcast it widely (e.g., via mass media), or let word spread from a small number of initially informed “seed” individuals? While conventional wisdom suggests broader dissemination is better, we show theoretically and experimentally that, once we take image concerns and related endogenous responses into account, this conclusion may be reversed. In a randomized field experiment during the chaotic 2016 Indian demonetization (Banerjee, Breza, Chandrasekhar, and Golub, 2019a), we varied how information about a change in the law was delivered to villages on two dimensions: how many were initially informed (broadcasting versus seeding) and whether the identity of the initially informed was publicly disclosed. Our results show that better learning outcomes can be achieved by giving fewer people information, as long as it is public that some people are informed and can be asked. The results are consistent with a model in which people need others’ help to make good use of announced information, but worry about signaling inability or unwillingness to comprehend the information they have access to. The stigma of information-seeking we identify can reinforce homophily in communication networks, leading to slow convergence in beliefs across groups and sustaining pockets of ignorance even when information is plentiful within a community.

### 3. FRAGILITY AND ROBUSTNESS IN FINANCIAL AND PRODUCTION NETWORKS

The last strand of research I discuss focuses on financial or production interdependencies modeled as networks. It considers shocks and the propagation of distress, their systemic consequences, and the implications of policies intended to foster resilience.

**3.1. Financial networks and contagion.** When firms experience defaults or shutdowns, value is lost not only by direct counterparties that have stakes in those firms (through debt, equity, or other claims), but also by indirect counterparties that have claims on those directly affected. The question of Elliott,

Golub, and Jackson (2014) is how the network of dependencies propagates the costs of shutdowns and how that ultimately redounds to the final claimants on an economy’s value.<sup>9</sup>

The model is a simple one in which institutions hold direct and indirect claims on each other, corresponding to a network structure. If an institution’s value crosses a critical threshold, a certain amount of value is lost by all its direct and indirect claimants, corresponding to a shutdown or default cost; such losses can cascade and trigger further discontinuous losses. The main results look at the expected losses as we vary the network structure governing the structure of claims. We show that the amount of damage caused by financial contagions can be nonmonotonic in both the diversification of the network (the typical number of direct counterparties) and in its integration (the magnitude of the typical financial relationship). Increasing either of these can exacerbate contagions but can also absorb shocks, and thus optimal policy responses can involve, for example, making a network more interconnected. Methodologically, the model combines simple Leontief computations of indirect interdependencies between different financial units with contagion models from random graph theory.

**3.2. Endogenous fragility in complex production.** The final project I discuss concerns networks in the real economy (as opposed to networks of financial interdependencies). The focus is the propagation of distress among firms that rely on each other for intermediate inputs. The systemic consequences of such risk have recently been a focus of work in macroeconomics and network theory (Acemoglu, Ozdaglar, and Tahbaz-Salehi, 2016; Baqaee, 2018). My current work shows that certain particularly stark systemic fragilities can arise because of the distinctive structure of production networks with failure-prone relational contracts.

Modern production is complex, featuring nested complementarities. For instance, the producer of a car needs to be able to source both dashboard electronics and engines. The production of each of these inputs, in turn, relies on being able to simultaneously source many complex components, each of which is crucial in the production of that input. In practice, these sourcing relationships are *specific*—they are formed between particular firms, rather than taking the form of anonymous transactions in a commodity market. Because these relationships are specific, they create risk: in a standard macroeconomic model, a firm relies on the total output of an industry, which has relatively low variance. When the firm relies on the output of specific counterparties, exposure to idiosyncratic risk is higher for each firm.<sup>10</sup>

In Elliott, Golub, and Leduc (2021) we extend supply network models to include these forces. We study a model in which the specific supply relationships just mentioned are used for sourcing, and in which firms can endogenously invest to insure against their failure. We find that the presence of specific sourcing relationships markedly changes the standard production network models at the aggregate level, and introduces new sorts of discontinuities. Even when individual firms optimally insure against supply link failure to maximize their own expected profits, the economy as a whole can be very sensitive to changes in aggregate parameters, such as the quality of contracting institutions. Small adjustments in aggregate productivity—which, in the standard model, have only a small effect on aggregate output—have very stark consequences in our model. For example, a fairly small negative shock (say, contracting becoming more difficult due to Brexit, or queues at ports due to Covid) can disrupt production and cause inventory to dry up suddenly across many seemingly unrelated sectors. *Simple* production, which does not require risky sourcing of multiple inputs, is not susceptible to this fragility.

We emphasize two main implications of our analysis. First, as economies develop, there can be sudden jumps in their ability to sustain complex production in equilibrium, and this occurs due to the structure of complex production networks. Second, the same phenomena create distinctive externalities. Counteracting seemingly small shocks to some industries can make a decisive difference for aggregate production possibilities, and this can be important as policymakers regulate and subsidize supply chain robustness.

<sup>9</sup>For overviews of the extensive work on this topic—foundational models, empirical motivations, etc.—see Glasserman and Young (2016) and Cabrales et al. (2016).

<sup>10</sup>The importance of relational contracts in production has been emphasized by Fafchamps and Minten (1999), Antràs (2005) and Acemoglu, Antràs, and Helpman (2007). In different settings, related problems of optimization in the presence of disruption risk have been studied by Blume, Easley, Kleinberg, Kleinberg, and Tardos (2011), Erol (2018), Erol and Vohra (2018), and Brummitt, Huremović, Pin, Bonds, and Vega-Redondo (2017), and many others.

This paper makes several contributions to the theory of economic networks, both at a conceptual and technical level. First, we introduce percolation analysis (i.e., disabling some links at random) to an otherwise standard network model of complex production—with *complex* meaning that each firm must source multiple inputs through customized relationships. That leads to the fragilities emphasized above. Second, as a modeling contribution, we demonstrate the tractability of studying equilibrium investments in links (more precisely, investments in the probability that links are operational) in such a setting. By defining a suitable model with a continuous investment choice and a continuum of nodes, investment problems are characterized by relatively tractable first-order conditions, because firms are able to average over the randomness in network realizations.<sup>11</sup> This technique is related to methods developed in Golub and Livne (2012), as well as in my student’s dissertation (Dasaratha, 2021), and we expect the methods to have other applications. Finally, using our equilibrium conditions to deduce the fragility results discussed above requires developing some new techniques to analyze the outcomes large network formation games. Outcomes depend in a subtle way on the structure of random graphs and incentives, reminiscent in some ways of forces studied by Jackson and Yariv (2007). This combination yields a complete analysis of equilibrium, which is what we need for our theory of aggregate fragility.

## REFERENCES

- ACEMOGLU, D., P. ANTRÀS, AND E. HELPMAN (2007): “Contracts and Technology Adoption,” *American Economic Review*, 97, 916–943.
- ACEMOGLU, D., K. BIMPIKIS, AND A. OZDAGLAR (2014): “Dynamics of Information Exchange in Endogenous Social Networks,” *Theoretical Economics*, 9, 41–97.
- ACEMOGLU, D., M. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): “Bayesian Learning in Social Networks,” *The Review of Economic Studies*, 78, 1201–1236.
- ACEMOGLU, D., C. GARCÍA-JIMENO, AND J. A. ROBINSON (2015): “State Capacity and Economic Development: A Network Approach,” *American Economic Review*, 105, 2364–2409.
- ACEMOGLU, D., A. OZDAGLAR, AND A. PARANDEHGHEIBI (2010): “Spread of (Mis)information in Social Networks,” *Games and Economic Behavior*, 70, 194–227.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2016): “Networks, Shocks, and Systemic Risk,” in *The Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoull’e, A. Galeotti, and B. W. Rogers, Oxford.
- ANGELETOS, G.-M. AND A. PAVAN (2007): “Efficient Use of Information and Social Value of Information,” *Econometrica*, 75, 1103–1142.
- ANTRÀS, P. (2005): “Incomplete Contracts and the Product Cycle,” *American Economic Review*, 95, 1054–1073.
- BALA, V. AND S. GOYAL (1998): “Learning from Neighbours,” *The Review of Economic Studies*, 65, 595–621.
- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): “Who’s Who in Networks. Wanted: The Key Player,” *Econometrica*, 74, 1403–1417.
- BANERJEE, A. (1992): “A Simple Model of Herd Behavior,” *The Quarterly Journal of Economics*, 107, 797–817.
- BANERJEE, A., E. BREZA, A. G. CHANDRASEKHAR, AND B. GOLUB (2019a): “When Less is More: Experimental Evidence on Information Delivery During India’s Demonetization,” NBER Working Paper No. 24679, available at SSRN 3163930; revise & resubmit at *The Review of Economic Studies*.
- BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2019b): “Using Gossips to Spread Information: Theory and Evidence from two Randomized Controlled Trials,” *The Review of Economic Studies*, 86, 2453–2490.
- BANERJEE, A. AND D. FUDENBERG (2004): “Word-of-Mouth Learning,” *Games and Economic Behavior*, 46, 1–22.
- BAQAEE, D. R. (2018): “Cascading Failures in Production Networks,” *Econometrica*, 86, 1819–1838.

<sup>11</sup>Rather than a large set of combinatorial inequalities, as in canonical network formation models.

- BEAMAN, L. AND J. MAGRUDER (2012): “Who Gets the Job Referral? Evidence from a Social Networks Experiment,” *American Economic Review*, 102, 3574–93.
- BENYISHAY, A. AND A. M. MOBARAK (2019): “Social Learning and Incentives for Experimentation and Communication,” *The Review of Economic Studies*, 86, 976–1009.
- BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2015): “Information and Volatility,” *Journal of Economic Theory*, 158, 427–465.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades,” *The Journal of Political Economy*, 100, 992–1026.
- BLUME, L., D. EASLEY, J. KLEINBERG, R. KLEINBERG, AND É. TARDOS (2011): “Which Networks are Least Susceptible to Cascading Failures?” in *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, IEEE, 393–402.
- BLUME, L. E., W. A. BROCK, S. N. DURLAUF, AND R. JAYARAMAN (2015): “Linear Social Interactions Models,” *Journal of Political Economy*, 123, 444–496.
- BRAMOULLÉ, Y., H. DJEBBARI, AND B. FORTIN (2009): “Identification of Peer Effects through Social Networks,” *Journal of Econometrics*, 150, 41–55.
- BRAMOULLÉ, Y., R. KRANTON, AND M. D’AMOURS (2014): “Strategic Interaction and Networks,” *American Economic Review*, 104, 898–930.
- BRUMMITT, C. D., K. HUREMOVIĆ, P. PIN, M. H. BONDS, AND F. VEGA-REDONDO (2017): “Contagious Disruptions and Complexity Traps in Economic Development,” *Nature Human Behaviour*, 1, 665.
- CABRALES, A., D. GALE, , AND P. GOTTARDI (2016): “Financial Contagion in Networks,” in *The Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. W. Rogers, Oxford.
- CALVÓ-ARMENGOL, A. AND M. O. JACKSON (2007): “Networks in Labor Markets: Wage and Employment Dynamics and Inequality,” *Journal of Economic Theory*, 132, 27–46.
- CALVÓ-ARMENGOL, A., J. MARTÍ, AND A. PRAT (2015): “Communication and Influence,” *Theoretical Economics*, 10, 649–690.
- CALVÓ-ARMENGOL, A., E. PATACCHINI, AND Y. ZENOU (2009): “Peer Effects and Social Networks in Education,” *The Review of Economic Studies*, 76, 1239–1267.
- CHANDRASEKHAR, A. G., B. GOLUB, AND H. YANG (2019): “Signaling, Shame, and Silence in Social Learning,” Working paper, Stanford University, available at <http://web.stanford.edu/~arungc/CGY.pdf>.
- CHANDRASEKHAR, A. G., H. LARREGUY, AND J. P. XANDRI (2020): “Testing Models of Social Learning on Networks: Evidence from two Experiments,” *Econometrica*, 88, 1–32.
- CHUNG, F., L. LU, AND V. VU (2004): “The Spectra of Random Graphs with Given Expected Degrees,” *Internet Mathematics*, 1, 257–275.
- DASARATHA, K. (2021): “Innovation and Strategic Network Formation,” *Review of Economic Studies*, forthcoming, arXiv:1911.06872.
- DASARATHA, K., B. GOLUB, AND N. HAK (2020): “Learning from Neighbors about a Changing State,” Available at SSRN 3097505; preliminary version appeared in EC ’18: Proceedings of the 2018 ACM Conference on Economics and Computation.
- DE MARTÍ, J. AND Y. ZENOU (2015): “Network Games with Incomplete Information,” *Journal of Mathematical Economics*, 61, 221–240.
- DEGROOT, M. H. (1974): “Reaching a Consensus,” *Journal of the American Statistical Association*, 69, 118–121.
- DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *The Quarterly Journal of Economics*, 118, 909–968.
- DUFLO, E. AND E. SAEZ (2003): “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment,” *The Quarterly Journal of Economics*, 118, 815–842.
- ELLIOTT, M. AND B. GOLUB (2019): “A Network Approach to Public Goods,” *Journal of Political Economy*, 127.

- ELLIOTT, M., B. GOLUB, AND M. O. JACKSON (2014): “Financial Networks and Contagion,” *American Economic Review*, 104, 3115–53.
- ELLIOTT, M., B. GOLUB, AND M. V. LEDUC (2021): “Supply Network Formation and Fragility,” Working paper, available at SSRN 3525459.
- ELLISON, G. AND D. FUDENBERG (1993): “Rules of Thumb for Social Learning,” *Journal of Political Economy*, 101, 612–643.
- (1995): “Word-of-mouth Communication and Social Learning,” *The Quarterly Journal of Economics*, 110, 93–126.
- EROL, S. (2018): “Network Hazard and Bailouts,” *Available at SSRN 3034406*.
- EROL, S. AND R. VOHRA (2018): “Network Formation and Systemic Risk,” *Available at SSRN 2546310*.
- EYSTER, E. AND M. RABIN (2010): “Naive Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–43.
- (2014): “Extensive Imitation is Irrational and Harmful,” *The Quarterly Journal of Economics*, 129, 1861–1898.
- FAFCHAMPS, M. AND B. MINTEN (1999): “Relationships and Traders in Madagascar,” *The Journal of Development Studies*, 35, 1–35.
- GALEOTTI, A., B. GOLUB, AND S. GOYAL (2020): “Targeting Interventions in Networks,” .
- GALEOTTI, A. AND S. GOYAL (2010): “The Law of the Few,” *American Economic Review*, 100, 1468–92.
- GLAESER, E. L., B. SACERDOTE, AND J. A. SCHEINKMAN (1996): “Crime and Social Interactions,” *The Quarterly Journal of Economics*, 111, 507–548.
- GLASSERMAN, P. AND H. P. YOUNG (2016): “Contagion in Financial Networks,” *Journal of Economic Literature*, 54, 779–831.
- GOLUB, B. AND M. O. JACKSON (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–149.
- (2012a): “Does Homophily Predict Consensus Times? Testing a Model of Network Structure via a Dynamic Process,” *Review of Network Economics*, 11.
- (2012b): “How Homophily Affects the Speed of Learning and Best-Response Dynamics,” *The Quarterly Journal of Economics*, 127, 1287–1338.
- GOLUB, B. AND Y. LIVNE (2012): “Strategic Random Networks: Why Social Networking Technology Matters,” Working paper, available at <http://www.people.fas.harvard.edu/~bgolub/papers/formation.pdf>.
- GOLUB, B. AND S. MORRIS (2020a): “Expectations, Networks and Conventions,” Available at SSRN 2979086.
- (2020b): “Higher-Order Expectations,” Available at SSRN 2979089.
- GOLUB, B. AND E. SADLER (2016): “Learning in Social Networks,” in *The Oxford Handbook of the Economics of Networks*.
- GOYAL, S. AND J. L. MORAGA-GONZALEZ (2001): “R&D Networks,” *Rand Journal of Economics*, 686–707.
- HELLWIG, C. AND L. VELDKAMP (2009): “Knowing What Others Know: Coordination Motives in Information Acquisition,” *The Review of Economic Studies*, 76, 223–251.
- JACKSON, M. AND L. YARIV (2007): “Diffusion of Behavior and Equilibrium Properties in Network Games,” *American Economic Review*, 97, 92–98.
- JACKSON, M. O. (2021): “Inequality’s Economic and Social Roots: The Role of Social Networks and Homophily,” *Available at SSRN 3795626*.
- LEVIN, D. A., Y. PERES, AND E. L. WILMER (2009): *Markov Chains and Mixing Times*, Providence, RI: American Mathematical Society, with a chapter on coupling from the past by James G. Propp and David B. Wilson.
- LEVY, G. AND R. RAZIN (2019): “Echo Chambers and their Effects on Economic and Political Outcomes,” *Annual Review of Economics*, 11, 303–328.
- MOBIUS, M., T. PHAN, AND A. SZEIDL (2015): “Treasure Hunt: Social Learning in the Field,” Tech. rep., National Bureau of Economic Research, working paper 21014.

- MOLAVI, P., A. TAHBAZ-SALEHI, AND A. JADBABAIE (2018): “A Theory of non-Bayesian Social Learning,” *Econometrica*, 86, 445–490.
- MOSSEL, E., A. SLY, AND O. TAMUZ (2015): “Strategic Learning and the Topology of Social Networks,” *Econometrica*, 83, 1755–1794.
- NIEHAUS, P. (2011): “Filtered Social Learning,” *Journal of Political Economy*, 119, 686–720.
- PENTA, A. (2011): “Multilateral Bargaining and Walrasian Equilibrium,” *Journal of Mathematical Economics*, 47, 417–424.
- SAMET, D. (1998): “Iterated Expectations and Common Priors,” *Games and Economic Behavior*, 24, 131–141.
- SHAHRAMPOUR, S., S. RAKHLIN, AND A. JADBABAIE (2013): “Online Learning of Dynamic Parameters in Social Networks,” *Advances in Neural Information Processing Systems*, 2013–2021.
- SPIELMAN, D. A. (2007): “Spectral Graph Theory and its Applications,” in *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)*, IEEE, 29–38.
- TAHBAZ-SALEHI, A. AND A. JADBABAIE (2009): “Consensus over Ergodic Stationary Graph Processes,” *IEEE Transactions on Automatic Control*, 55, 225–230.
- VIVES, X. (1993): “How fast do rational agents learn?” *The Review of Economic Studies*, 60, 329–347.
- YILDIZ, M. (2003): “Walrasian Bargaining,” *Games and Economic Behavior*, 45, 465–487.

NORTHWESTERN UNIVERSITY, BGOLUB@NORTHWESTERN.EDU.