OVERVIEW OF RESEARCH

BENJAMIN GOLUB

Abstract. This document is an overview of my research. It is organized (loosely) into several themes: social learning; externalities and negotiations; strategic spillovers and contagion; strategic network formation; and empirical work on obstructions to information transmission.

My research focuses on social and economic networks. A recurring theme is describing economically important aspects of networks through theory-based summary statistics that can be useful in empirical studies and policy analyses.

1. Social learning and information aggregation

Social learning and influence are central to consumer choice, financial decisions, and political behavior. One strand of my research considers agents in a social network being influenced by each other over time. The updating of their beliefs and behaviors is modeled by simple "DeGroot" rules in which the past opinions of one's contacts are aggregated by averaging them.

1.1. **Naive learning, influence, and the wisdom of crowds.** Suppose that, at the beginning of time, nature draws a state θ from, say, a standard normal distribution. Agents start out with noisy estimates of this state based on private information, and then they repeatedly learn from their network neighbors using the DeGroot rule described above. After sufficient time, under some assumptions on network structure, agents will converge to a consensus belief. [13] asks which features of networks ensure that large groups will learn well, even if their updating rules are naive. There is a simple, network-based measure of influence that determines how each individual's idiosyncratic noise enters the group's eventual consensus. Convergence to truth obtains if and only if the influence of the most influential agent in the society is vanishing as the society grows. Small groups that are prominent or central—influencing many others directly or indirectly—turn out to be an important obstacle to efficient learning.

Date: May 27, 2018.

1.2. **Homophily and the speed of convergence.** This strand of research continues with a series of papers [16, 15, 17] focused on how social learning is affected by homophily, the tendency of people to socialize most with demographically similar others.[1] Broadly speaking, homophily slows the process of converging to a consensus. In making this point, the study of learning dynamics is helped by the nice algebraic structure of DeGroot updating processes, which is closely related to the dynamics of Markov chains. [17] uses this algebraic structure to derive new bounds on the speed of convergence of learning to equilibrium in terms of homophily. [16] combines this structure with techniques from random graph theory to show that the speed of convergence can be described in terms of a group-level measure of segregation called *spectral homophily*. Here, the techniques from probability help identify what features of large-scale network structure are essential to the speed of learning, and which ones can be ignored as noise.

1.3. **Statistics of networks and diffusion processes.** The characterization just discussed is reliant on suitable statistical assumptions about the underlying network. In [15] we find that these statistical assumptions seem to hold up in the Add Health study of high school friendship networks. For the purpose of predicting the behavior of DeGroot processes of learning and information aggregation, we can restrict attention to fairly coarse measures of network structure.

[14] examines different statistical issues that are relevant in the empirical estimation of information diffusion—sampling biases that naturally arise in our observation of these processes. After accounting for these biases, simple branching process models turn out to be a remarkably good fit to viral diffusion on the Internet.

1.4. **A changing environment and a Bayesian approach to linear learning rules.** Consider now a group learning, over a period of time, about an *evolving* fundamental state, such as the future value of an asset [4]. The state, θ_t, drifts around according to a stationary, discrete-time AR(1) process given by $\theta_{t+1} = \rho\theta_t + \nu_{t+1}$, and agents receive conditionally independent Gaussian signals of its current value. In each period, each agent observes an independent, normal signal $s_{i,t} \sim \mathcal{N}(\theta_t, \sigma_i^2)$ of the current state andsome past estimates of her neighbors in an arbitrary network, which (because they depend on recent states) are relevant for estimating θ_t. The agent's action is an estimate $a_{i,t}$: she sets it to her expectation, given all her information, of the current state θ_t. Her estimate is then used by her neighbors in the next round in the same way. When agents are Bayesian, stationary equilibrium learning rules take a simple, time-invariant form: agents form their

[1] The results of [13] and [16] formed two chapters of my Ph.D. dissertation [12]. The remaining chapter is an early version of [7], discussed below.

next-period estimates by taking linear combinations of each other's earlier estimates and their own private signals. This provides a Bayesian microfoundation for the form of DeGroot learning rules.

Substantively, we ask: when can a group aggregate information quickly, keeping up with the changing environment? First, if private signal distributions are diverse enough across agents, then Bayesian learning achieves good information aggregation as long as individuals observe sufficiently many others. Second, without such diversity, Bayesian information aggregation can fall far short of good aggregation benchmarks, and can be Pareto-inefficient. Third, good aggregation requires anti-imitation (placing negative weight on some neighbors); when we consider linear rules where the weights are not optimal and do not have this feature, agents' estimates are inefficiently confounded by "echoes."

1.5. **Perverse outcomes of expert advice.** Another project about information aggregation studies forces that can make decisions worse even as the experts who advise decisions become better able to discern the best decision [9]. We model two experts who must make predictions about whether an event will occur or not. The experts receive private signals about the likelihood of the event occurring, and simultaneously make one of a finite set of possible predictions, corresponding to varying degrees of alarm. The information structure is commonly known among the experts and the recipients of the advice. Each expert's payoff depends on whether the event occurs and her prediction. Our main result shows that when either or both experts receive uniformly more informative signals, for example by sharing their information, their predictions can become unambiguously less informative. We call such information improvements perverse. Suppose a third party wishes to use the experts' recommendations to decide whether to take some costly preemptive action to mitigate a possible bad event. Regardless of how this third party trades off the costs of various errors, he will be worse off after a perverse information improvement.

1.6. **A Survey.** [23] surveys progress in social learning theory—both in the information cascades/herding paradigm and in the DeGroot models discussed above.

2. A network perspective on externalities and negotiations

A second agenda uses networks to analyze payoff externalities and strategic spillovers. [7] study how groups cooperate when externalities are complex. For example, countries can, at a cost, pollute less. Due to geography and other asymmetries, the benefits of a given such effort are not distributed uniformly. This paper the prospects for cooperation in such a setting by defining a network reflecting marginal externalities.[2] The spectral radius of the externality network quantifies the collective returns to investment in public goods. We use this measure to characterize players who are essential to negotiations, and to describe when negotiations can be subdivided without much loss. Methodologically, the results open new connections between the structure of general equilibrium, on the one hand, and the theory of networks and centrality measures, on the other. [6] presents axiomatic measures for the spectral measure of scope for negotiations that we study in [7].

Within a firm, an important type of externality occurs due due to delays in processing. [20] uses a network approach to analyze how to manage these externalities. We find conditions on the structure of production that allow simple transfer pricing schemes to achieve efficient outcomes.

3. Strategic spillovers and contagion in networks

The next set of projects looks at strategic linkages among a network of interacting units and how that affects the stability and efficiency of economic systems.

3.1. **Financial networks.** When firms experience defaults or shutdowns, value is lost 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 [8] is how the network of dependencies propagates the costs of shutdowns and how that ultimately redounds to consumers. 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 stakes of the typical financial relationship). Increasing either of these can exacerbate contagions but can also absorb shocks. The paper suggests ways for policymakers to assess the tradeoff. A related project, [1], studies externalities across two different financial networks, e.g., over-the-counter markets for two different kinds of financial claims. We posit that firms' insolvency spills over across networks, and show that this can make the contagions of [8] more stark and discontinuous.

3.2. **Network games of incomplete information.** This agenda focuses on games in networks under incomplete information [21, 22]. Individuals interact strategically with their network neighbors. For instance, when team members choose their level of effort, incentives are affected by their own attributes (ability, etc.) and the efforts exerted by

[2]These results were first announced in [5].

their collaborators. When firms invest in capacity, they optimize in response to costs and benefits idiosyncratic to their own business, as well as other firms' investment decisions.

We model agents who differ in (i) whom their payoffs depend on; (ii) what information they have; and (iii) how they interpret information (their priors). Our goal in [21] is to characterize how equilibrium behavior and welfare depends on these dimensions.

The key methodological innovation is a new representation, called the *interaction structure*, that captures network relationships and information asymmetries simultaneously. Using this representation along with Markov chain techniques, we unify existing results from network games and classical incomplete-information games, such as beauty contests. We also characterize outcomes under assumptions that go beyond standard ones—e.g., relaxing the common prior assumption.

We illustrate the techniques by studying their consequences in several applications. For example, we ask which member of a team is the most influential, in the sense of their private preferences making the biggest difference to equilibrium outcomes. Under certain conditions, the most influential agent will be the one who is most central in the collaboration network, but under different conditions (e.g., without common priors) it can be the one who has the worst private information about the state.

In [22] we show that the characterizations above are related to fundamental questions about the structure of iterated expectations. This paper generalizes known results about the connection between iterated expectations and common priors, and introduces new techniques for the study of the connection between these objects.

3.3. **Targeting interventions in networks.** When strategic spillovers are present, a natural question is how a planner should shape incentives (e.g., individual students' rewards from studying, or individual firms costs) in pursuit of an aggregate goal, such as maximizing welfare or minimizing volatility. In [11], we analyze a variety of targeting problems by identifying how a given profile of incentive changes is amplified or attenuated by the strategic spillovers in the network.

We introduce a method of decomposing any potential intervention into *principal components* determined by the network. A particular ordering of principal components describes the planner's priorities across a range of network intervention problems. Optimal policies are simplest when the budget for intervention is large. 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. In games of strategic substitutes, the optimal intervention is very different: it moves neighbors' incentives in opposite directions, dividing local communities into positively and negatively targeted agents, with few links across these two categories.

3.4. **Linking groups and the leverage of weak ties.** In [11], the optimal targeting policy when actions of neighbors are strategic complements is to target individuals in proportion to their *eigenvector centrality* (also known as PageRank). In [18] we focus on this measure, which often comes up in networks, and characterize how the centrality of each member of a society changes when initially disconnected groups begin interacting with each other via a new bridging link. Arbitrarily weak intergroup connections can have arbitrarily large effects on the distribution of centrality. For instance, if a high-centrality member of one group begins interacting symmetrically with a low-centrality member of another, the latter group has the larger centrality in the combined network in inverse proportion to the centrality of its "emissary." We also find that agents who form the intergroup link, the "bridge agents," become relatively more central within their own groups, while other intragroup centrality ratios remain unchanged. Thus, features of strategic outcomes can hinge in very sensitive ways on fairly "small" details of the link structure, provided those small changes occur at the right, high-leverage points.

4. Strategic network formation and network fragility

4.1. **A search model of socializing.** One of the distinctive aspects of the economic approach to networks is to view them as a kind of capital in which individuals rationally invest. However, the standard equilibrium concepts used to study network formation, such as pairwise stability, are very demanding: agents carefully consider each link in the context of the whole network. In [19] we develop an alternative model inspired by search theory: agents rationally choose an overall socializing or search effort, anticipating value from both direct and indirect links. Chance meetings then result in a complex random network. Random graph theory permits precise predictions about welfare, despite the presence of multiple equilibria. The main substantive predictions are about sudden transitions from efficient to inefficient socializing: small changes in socializing costs can lead to large drops in aggregate social capital.

4.2. **Endogenous fragility in complex production.** Similar fundamental forces apply in a quite different setting: that of modern supply chains [10]. Modern production is complex, and whether it is organized within a firm's boundaries or through supply networks involving many firms, relationships are crucial. Successful production requires many relationships to be functioning well at once, but there are many reasons why supply relationships fail and collaborations within a firm can be unproductive. While operating in a suitable environment with good institutions can mitigate these risks, individuals make contributions towards such an environment that matter, and there is the potential for extensive free-riding in these contributions. How then can the right environment be maintained in the presence of decentralized interactions among agents who are numerous and relatively short-lived? We present a theory of how incentives endogenously arise to contribute to the right conditions for complex production. The theory implies that those conditions, and the complex production they foster, are necessarily fragile.

5. Stigma in social learning: Theory, evidence, and policy

This program focuses on how agents decide whether to participate in information exchange, whom to talk to, and what to ask. Social learning is important in a variety of settings. Practical guidance from peers about education and job opportunities can be life-changing. Social learning can also counteract deceptive messages, such as financial scams or fake news. In many cases, however, asking questions may be stigmatized. This can be explained in a signaling model where people are instrumentally concerned about making others think well of them. It can also come from a psychological phenomenon distinct from signaling, which we call shame, in which people simply dislike seeking advice when they are thought to be of low ability. In this agenda, we examine these forces, we study whether the information-sharing processes that are the focus of the models in Section 1 get activated in the first place

5.1. **A experiment on stigma in seeking information.** In [3], we develop a model where, due to signaling or shame deterrents, low-ability individuals may not seek the information they need. To measure the extent to which stigma operates in practice, we analyze an experiment across 70 villages in India. We find that low-ability individuals do face large stigma effects (reducing their seeking of information by a factor of two). Both signaling and shame operate, and which force is prevalent depends on the pre-existing social relationships between the people involved.

5.2. **Policy implications and an experimental information to combat misinformation.** We then 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 stigma and other "higher-order" concerns into account, this may not be the case.

In a randomized field experiment during the chaotic 2016 Indian demonetization [2], 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, provided people also come to believe that they can learn what they don't know by asking others. 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 and sustaining pockets of ignorance even when information is plentiful within a community.

5.3. **Stigma and brokerage.** Of course, stigma and mistrust operate beyond the domain of information-sharing alone. The institution of brokerage, whereby individuals transact across social boundaries, is subject to many frictions and forms of distrust, including stigma. In [24], we use sociological and economic models to explain the ways in which brokerage is a fragile relationship.We also examine social mechanisms that can stabilize brokerage relationships. Each of these mechanisms rests on the emergence or existence of supporting institutions. We suggest that the "grafting" of brokerage functions onto existing institutions, such as the provision of public goods, may be the most stable and effective resolution to the tensions inherent in brokerage, but it is also the most institutionally demanding.

References

[1] Aymanns, C., C.-P. Georg, and B. Golub (2018): "Illiquidity Spirals in Coupled Over-the-Counter Markets," University of St. Gallen, School of Finance Research Paper No. 2018/10. Available at SSRN: https://ssrn.com/abstract=2881814.

[2] Banerjee, A., E. Breza, A. G. Chandrasekhar, and B. Golub (2018): "When Less is More: Experimental Evidence on Information Delivery during India's Demonetization," .

[3] Chandrasekhar, A. G., B. Golub, and H. Yang (2017): "Signaling, Shame, and Silence in Social Learning," Working paper, Stanford University, available at http://web.stanford.edu/∼arungc/CGY.pdf.

[4] Dasaratha, K., B. Golub, and Hak (2018): "Social Learning in a Dynamic Environment," Available at SSRN: https://ssrn.com/abstract=3097505.

[5] Elliott, M. and B. Golub (2013): "A network approach to public goods," in *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, ACM, 377–378.

[6] Elliott, M. and B. Golub (2015): "Ranking Agendas for Negotiations," Available at http://www.people.fas.harvard.edu/∼bgolub/papers/agendas.pdf.

[7] Elliott, M. and B. Golub (2018): "A network approach to public goods," *Journal of Political Economy*, Forthcoming.

[8] Elliott, M., B. Golub, and M. O. Jackson (2014): "Financial networks and contagion," *American Economic Review*, 104, 3115–53.

[9] Elliott, M., B. Golub, and A. Kirilenko (2014): "How Sharing Information Can Garble Experts' Advice," *American Economic Review (Papers and Proceedings)*, 104, 463–68.

[10] Elliott, M., B. Golub, and M. V. Leduc (2018): "Endogenous fragility in complex economic systems," Working paper.

[11] Galeotti, A., B. Golub, and S. Goyal (2017): "Targeting Interventions in Networks," Available at SSRN: https://ssrn.com/abstract=3054353.

[12] Golub, B. (2011): "Essays on Economic Networks," Ph.D. thesis, Stanford University.

[13] Golub, B. and M. Jackson (2010): "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2, 112–149.

[14] Golub, B. and M. O. Jackson (2010): "Using selection bias to explain the observed structure of internet diffusions," *Proceedings of the National Academy of Sciences*, 107, 10833–10836.

[15] Golub, B. and M. O. Jackson (2012): "Does homophily predict consensus times? Testing a model of network structure via a dynamic process," *Review of Network Economics*, 11.

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

[17] Golub, B. and M. O. Jackson (2012): "Network structure and the speed of learning measuring homophily based on its consequences," *Annals of Economics and Statistics*, 33–48.

[18] Golub, B. and C. Lever (2010): "The leverage of weak ties: How linking groups affects inequality," Available at http://www.people.fas.harvard.edu/∼bgolub/papers/intergroup.pdf.

[19] 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.

[20] Golub, B. and R. P. McAfee (2011): "Firms, queues, and coffee breaks: a flow model of corporate activity with delays," *Review of Economic Design*, 15, 59–89.

[21] Golub, B. and S. Morris (2017): "Expectations, Networks and Conventions," Available at SSRN: https://ssrn.com/abstract=2979086.

[22] Golub, B. and S. Morris (2017): "Higher-Order Expectations," Available at SSRN: http://ssrn.com/abstract=2979089.

[23] Golub, B. and E. Sadler (2016): "Learning in Social Networks," in *The Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, B. Rogers, and B. Rogers, Oxford University Press, chap. 19, 504–542.

[24] Stovel, K., B. Golub, and E. M. M. Milgrom (2011): "Stabilizing brokerage," *Proceedings of the National Academy of Sciences*, 108, 21326–21332.

Department of Economics, Harvard University, Cambridge, U.S.A., bgolub@fas.harvard.edu.